

A Peered Bulletin Board for Robust Use in Verifiable Voting Systems

Chris Culnane and Steve Schneider
Department of Computing, University of Surrey

January 17, 2014

Abstract

The Web Bulletin Board (WBB) is a key component of verifiable election systems. It is used in the context of election verification to publish evidence of voting and tallying that voters and officials can check, and where challenges can be launched in the event of malfeasance. In practice, the election authority has responsibility for implementing the web bulletin board correctly and reliably, and will wish to ensure that it behaves correctly even in the presence of failures and attacks. To ensure robustness, an implementation will typically use a number of peers to be able to provide a correct service even when some peers go down or behave dishonestly. In this paper we propose a new protocol to implement such a Web Bulletin Board, motivated by the needs of the vVote verifiable voting system. Using a distributed algorithm increases the complexity of the protocol and requires careful reasoning in order to establish correctness. Here we use the Event-B modelling and refinement approach to establish correctness of the peered design against an idealised specification of the bulletin board behaviour. In particular we show that for n peers, a threshold of $t > 2n/3$ peers behaving correctly is sufficient to ensure correct behaviour of the bulletin board distributed design. The algorithm also behaves correctly even if honest or dishonest peers temporarily drop out of the protocol and then return. The verification approach also establishes that the protocols used within the bulletin board do not interfere with each other. This is the first time a peered web bulletin board suite of protocols has been formally verified.

1 Introduction

Verifiable voting systems such as Prêt à Voter [CRS05, RBH⁺09], Scantegrity [CCC⁺10], Helios[Adi08], Wombat [BNFL⁺12], STAR-Vote [BBB⁺13] and Civitas [CCM08] typically have a requirement to publish information concerning votes cast and how they have been processed, in order to provide verifiability. Voters and other external parties are able to check the published information and challenge the election if any cheating has occurred. Such systems are generally described using a “Bulletin Board” for publication: a repository of the information collected throughout the election, made publicly available for inspection.

There are certain (generally implicit) security assumptions on the bulletin board: that once items are on the bulletin board then they will not be removed, that the final information

given at the end of the election is fixed and cannot be adjusted, and that it will provide the same view of that information to all parties. For example, Adida’s characterisation [Adi06] states that “Cryptographic voting protocols revolve around a central, digital bulletin board. As its name implies, the bulletin board is public and visible to all, via, for example, phone and web interfaces. All messages posted to the bulletin board are authenticated, and it is assumed that any data written to the bulletin board cannot be erased or tampered with.” Alternatively a bulletin board has been described as a “broadcast channel with memory” [Pet05, CGS97, KTV12], with a Web Bulletin Board treated as a public broadcast channel.

Achieving these properties in an implementation is not so straightforward. A current view is that “*we don’t know how to build a secure bulletin board*” [Wag13], and to date there is no generally available implementation of a secure bulletin board. In practice bulletin boards are generally implemented by collecting election information as it progresses, and publishing the information via a website, as done for example by Helios, Wombat, and STAR-Vote, or making it available via a git repository as in the Norway 2013 e-voting trial [Nor13]. However, these are not tamper proof, and information can be changed on them unless there are additional safeguards such as the cryptographic mechanisms based on hash chains proposed by Heather and Lundin [HL08]. The design of STAR-VOTE uses multiple peers to tolerate faulty or malicious components, and has the election authority sign the bulletin board contents, thus changes can occur only with the collusion of the electoral authority.

The bulletin board presented in this paper arises from the need to implement a bulletin board as part of the vVote system being developed for the Victorian State election 2014 [BCH⁺12]. The Victorian State election runs over a two week period of “early voting” before election day itself, and the bulletin board is required to publish its information daily during the election. For robustness and trust the bulletin board will be comprised of a number of peers to receive items, provide receipts, and publish information. The rate at which votes may be received means that the peers cannot sustain the overhead of a consensus protocol every time an item is posted, so they each maintain a local copy of their view of the bulletin board, and agree on the bulletin board only when it is time to publish. A further challenge is that the bulletin board may need to reject some items, for example audit of a ballot previously used to vote, or any vote on a ballot previously used or audited, so that incompatible posts are not published. We achieve this requirement provided a threshold of the peers are honest and operational the bulletin board will behave correctly, even in the presence of individual peers going down, external attacks and a minority of dishonest peers.

This paper presents a new bulletin board protocol designed to run with a network of peers and to operate correctly when a threshold of the peers are honest and operational. We provide a formal model and verification of the protocol, using the framework of Event-B [Abr10]. We verify the protocol in the context of a Dolev-Yao attacker [DY83], who has control over the network and a minority of peers.

The paper is structured as follows: Section 2 presents and motivates the protocol, Section 3 introduces the Event-B framework and the refinement approach to modelling and verification, and the way it will be applied to the protocol and Sections 4–7 provide the details of the four stages of modelling aspects of the protocol and the verification proofs in terms of simulation. Section 8 discusses sufficient conditions for liveness, and Section 9 concludes with a discussion of what has been achieved, its relationship to related work, and its context.

2 A peered bulletin board protocol

We present an implementation of a bulletin board that accepts items to be posted (if they do not clash with previous posts), issues receipts, and periodically publishes what it has received. The bulletin board published for any particular period must include all items that had receipts issued during that period. Robustness is achieved through the use of several peered servers which cooperate on accepting items, issuing receipts, and publishing the bulletin board. They make use of a threshold signature scheme which allows a subset of the peers above a particular threshold to jointly generate signatures on data. The peers collectively provide the bulletin board service as long as a threshold of them are honest, and as long as a threshold of them are involved in handling any item posted to the bulletin board. Thus the implementation is correct in the presence of communication failures, unavailability or failure of peers, and also dishonesty of peers. The threshold t required to achieve this must be greater than two-thirds of the total number n of peers: $t > 2n/3$. There is no single point of failure: the system can tolerate failure or non-participation of any component, as long as a threshold of peers remain operational at any stage. It also allows for different threshold sets of peers to be operational at different times. For example, a peer may be rebooted during the protocol, missing some item posts, and may then resume participation.

The key properties we require for this bulletin board are:

- (bb.1) only items that have been posted to the bulletin board may appear on it;
- (bb.2) any item that has a receipt issued must appear on the published bulletin board;
- (bb.3) two clashing items must not both appear on the bulletin board;
- (bb.4) items cannot be removed from the bulletin board once they are published.

It follows from bb.2 and bb.3 that if two items clash then receipts must not be issued for both of them.

The bulletin board provides a protocol for the posting of an item and its acknowledgement with a receipt, and provides another two related protocols for the publishing of the bulletin board: an optimistic one, and a fallback.

2.1 Posting and acknowledgement

The protocol for posting an item x in period p , and issuing the acknowledgement, is as follows:

1. $User \rightarrow P_i : x$ (for each $i \in I$)
each P_i checks no clash between x and previous posts
2. $P_i \rightarrow P_j : sig_{sk_i}(p, x)$ (for each $i, j \in I, j \neq i$)
each P_i waits for at least a threshold number of signatures
3. $P_i \rightarrow User : sig_{ssk_i}(p, x)$ (for each $i \in I$)




Figure 1: Posting Protocol

To post an item x , the User should first send x to each of the peers, as shown in Round 1. Each peer checks that x does not clash with any posts it has received previously (from the current period or previous periods). The peers then sign (p, x) with their own individual signing key, and send the result to each of the other peers, as shown in Round 2. Peers store all of the received signatures into their local database. Finally, once a peer has obtained a threshold number of signatures on (p, x) (including its own), it sends its share of the threshold signature on (p, x) back to the User. Once the User has received a threshold number of such shares it is able to combine them to provide a signature on (p, x) , and this serves as the receipt. This protocol is shown in Figure 1. It is repeated for each item to be posted in the period.

2.2 Publishing the Bulletin Board

The bulletin board is published at the end of the period. The aim is for the peers to agree on the contents of the bulletin board and to issue their signature share on it to a public hosting service that can combine the signature shares and make the resulting signature publicly available.

Peer i 's local record of the bulletin board $B_{i,p}$ is those items that it has received a threshold number of signatures on, which are those items it issues a signature share on towards the




Figure 2: Optimistic Protocol

receipt.

The peers first of all run an optimistic protocol: this will succeed if at least a threshold of the local bulletin boards agree, which will be the case in practice if all peers are working properly. The optimistic protocol is given as follows:

1. $P_i \rightarrow P_j$: $sig_{sk_i}(p, h(B_{i,p}))$ (for each $i, j \in I, j \neq i$)
each P_i checks the hashes from a threshold of peers agree
2. $P_i \rightarrow WBB$: $B_{i,p}, sig_{ssk_i}(p, h(B_{i,p}))$ (for each $i \in I$)

The peers each sign a hash of their local copy of the bulletin board, and send them to each other. If a threshold agree then they can issue the bulletin board and a share of the threshold signature on the hash. This is illustrated in Figure 2.

If the optimistic protocol does not run successfully, because the hashes do not agree, that indicates that local bulletin boards are different. In this case the peers exchange information about their bulletin boards using the fallback protocol as follows:

1. $P_i \rightarrow P_j$: $D_{i,p}$ (for each $i, j \in I, j \neq i$)
each P_i adds any missing information received from others to its own database




Figure 3: Fallback Protocol

Each peer sends its database of signatures it has collected from the posting period to all the other peers, which update their databases with any signatures that are missing. They can then recalculate their local bulletin board. This is illustrated in Figure 3.

After the fallback protocol is completed, the peers return to the optimistic protocol and repeat. This is only required once for our liveness assumptions. We assume for liveness either (1) that all peers are online and able to communicate during the fallback protocol (with no assumptions about the posting phase or correct behaviour of users), or (2) that a threshold of honest peers are online and able to communicate during the fallback protocol, and at every stage of the posting phase a threshold set of peers were live and able to communicate and that the posting users behaved honestly. Under either of these two assumptions only one round of the fallback protocol is needed. The difference with Byzantine Agreement protocols, which tend to require up to $(n - t) + 1$ rounds to achieve agreement, is that the databases the peers start with have some consistency between them. If thresholds of peers received posts correctly in the posting phase, then the honest peers involved in the exchange of information in the fallback protocol will all obtain the full bulletin board after one round. Further explanation is provided in Section 8.




Figure 4: Template of an Event-B machine and an event.

3 Modelling and Verification Framework

We use the action systems approach of Event-B [Abr10, MAV05] as our formal framework to model the protocol and to verify it. In this approach systems are described in terms of the **states** that they can be in, and the **events** that transform the state.

A system is defined as a *machine*, which encapsulates its state, and its events. State information is described in terms of state variables and invariants on them. The machine describes how the state is initialised, and how it can be updated with events.

The Event-B approach supports *refinement*, a relationship showing when one system implements another. This approach allows a specification to be captured as an ideal machine that expresses the required behaviour. An implementation satisfies the specification if it is a refinement.

Figure 4 illustrates how a machine is defined. Machine M is given with a list of state variables v , a state invariant $I(v)$, and a set of events ev, \dots to update the state. Initialisation is a special event *init* which sets the initial state of the machine, and its guard is *true*.

Each event has a *guard* $G(v)$ over the variables v , and a *body*, usually written as an assignment S on the variables. The assignment is associated with a *before-after predicate* $BA(v, v')$ describing changes of variables upon event execution, in terms of the relationship between the variable values before (v) and after (v'). For example, the assignment $v := v + 1$ is associated with the predicate $v' = v + 1$. The body can also be written as $v :| BA(v, v')$, which assigns to v any value v' which makes the predicate $BA(v, v')$ true (see right of Fig. 4, where BA is the predicate in event *evt*). In Event-B an event may also introduce local variables, which can be included in the guard (which constrains what values they can take), and in the body where they can be used to define the change of state. Such events are constructed as:

$$\begin{aligned}
 evt \hat{=} & \\
 & \mathbf{any} \ x \\
 & \mathbf{where} \ G(v, x) \\
 & \mathbf{then} \ v :| BA(v, x, v') \\
 & \mathbf{end}
 \end{aligned}$$

Some of the conditions on x may be included in the **any** clause rather than the **where** clause for readability (see e.g. *post* and *a_msg1* of Figure 7). Nondeterministic assignment has its own syntax: $x \in S$ assigns x some arbitrary element of S . This is an abbreviation for

any s **where** $s \in S$ **then** $x := s$ **end**.

In this paper, all events have some feasible final state: whenever $G(v, x)$ is true then there is some v' such that $BA(v, x, v')$ holds.

The Event-B approach to semantics, provided in [Abr10, MAV05], is to associate proof obligations with machines. The key proof obligation on an event is that it preserves the invariant: when an event is called within its guard, then the state resulting from executing the body should meet the invariant. For example, in the case of the machine in Fig. 4 we obtain the following proof obligation **INV** on events which have the form of *evt*. It states that if the invariant I holds on v , and the guard $G(v)$ is true, and the before-after predicate relates v' to v , then the invariant I should be true on the state v' reached after the event:

$\begin{array}{c} I(v) \wedge G(v) \wedge BA(v, v') \\ \vdash \\ I(v') \end{array}$	INV
---	------------

Discharging this proof obligation establishes that the event preserves the invariant. The machine is consistent if this is true for all of its events. It is true for all events in all machines presented in this paper: establishing this is one part of the proof of correctness.

3.1 Event-B refinement

In Event-B, the intended refinement relationship between machines is directly written into the refinement machine definitions. As a consequence of writing a refining machine, a number of proof obligations arise. Here, a machine and its refinement take the following form:

machine M_0 variables v invariant $I(v)$ events $init_0, ev_0, ev'_0, \dots$ end	machine M_1 refines M_0 variables w invariant $J(v, w)$ events $init_1, ev_1, ev'_1, \dots$ end
---	--

The machine M_0 is refined by machine M_1 , written $M_0 \preceq M_1$, if the given *linking invariant* $J(v, w)$ on the variables of the two machines is established by their initialisations, and preserved by all events. Any transition performed by a concrete event of M_1 can be matched by a step of the corresponding abstract event of M_0 , or matched by *skip* for newly introduced events, in order to maintain J . This is similar to the approach of downwards simulation data refinement [DB01], where the *simulation relation* plays the role of the linking invariant. Formally, the refinement relation $M_0 \preceq M_1$ between abstract machine M_0 and concrete machine M_1 holds if the following proof obligations given below hold for all events:

GRD_REF: Guard Strengthening If a concrete event matches an abstract one, then this rule requires that when the concrete event is enabled, then so is the matching abstract one. The rule is:

$\begin{array}{c} I(v) \wedge J(v, w) \wedge H(w) \\ \vdash \\ G(v) \end{array}$	GRD_REF
--	----------------

INV_REF: Simulation This ensures that the occurrence of events (including initialisation) in the concrete machine can be matched in the abstract one. If there is a matching abstract event then the rule is:

$\begin{array}{c} I(v) \wedge J(v, w) \wedge H(w) \wedge BA1(w, w') \\ \vdash \\ \exists v'.(BA0(v, v') \wedge J(v', w')) \end{array}$	INV_REF₁
--	----------------------------

New events are treated as refinements of *skip*. In this case the abstract state does not change (i.e., $v' = v$), and the rule is

$\begin{array}{c} I(v) \wedge J(v, w) \wedge H(w) \wedge BA1(w, w') \\ \vdash \\ J(v, w') \end{array}$	INV_REF₂
--	----------------------------

Refinement with respect to A

It may be that an environment interacts with a machine M_0 only on some subset A of its events. In that case we can consider a refinement M_1 of M_0 with respect to A . This requires that M_1 also has all the events A , and that **GRD_REF** and **INV_REF₁** must hold for all the events in A . However, other events of M_1 can be matched either by *skip*, or by some matching event (not in A) in M_0 , in which case the guard must also match. Thus for events not in A we weaken the requirement to the single proof obligation **GRD_INV_REF₃**:

$\begin{array}{c} I(v) \wedge J(v, w) \wedge H(w) \wedge BA1(w, w') \\ \vdash \\ J(v, w') \vee (G(v) \wedge \exists v'.(BA0(v, v') \wedge J(v', w'))) \end{array}$	GRD_INV_REF₃
--	--------------------------------

We will use this notion of refinement to express our requirements on the bulletin board protocol.

3.2 Framework for Bulletin Board Modelling and Verification

We are concerned with developing a peered bulletin board that can operate correctly in an unreliable environment, and with some potentially misbehaving peers. In particular, communications between the bulletin board and its users may be under the control of an




Figure 5: *BBSpec*: ideal bulletin board and communication medium

adversary, who may intercept, divert, block, duplicate and spoof messages. The bulletin board is designed for use in such an environment.

The specification of the bulletin board will encapsulate the required behaviour. This will be described as an Event-B model *BBSpec* with a description in terms of the architecture shown in Figure 5, of an ideal bulletin board in the context of a reliable communication medium. Users may use the events *post*, *ack* and *publish* to interact with the bulletin board, but communication occurs via the medium. The bulletin board has its own corresponding interactions with the medium, labelled *a_msg1*, *a_msg2* and *a_msg3*. These events are also within the model *BBSpec*, but they are not accessible directly to users. Hence it is the behaviour of *BBSpec* on the set of events $\{post, ack, publish\}$ that must be matched by any implementation.

The bulletin board implementation uses a number of peers, for robustness and in order to distribute trust. There are a total of n peers, and we use a threshold signature scheme in which we require t shares in order to produce a signature. Our model of the protocol will be an Event-B model *BBProt*, in which we consider the adversary to control the communication medium to and from the peers and the *WBB*, and between them. Hence any communication can be blocked. We also consider that the adversary can control up to $n - t$ peers. This means that such peers can sign and create any messages for sending, whether or not such messages are in accordance with the protocol, provided they have the appropriate keys.

We consider that (at least) a threshold t of the n peers are honest: that they follow the protocol. Without loss of generality we will consider peers 1 to t to be honest, and $t + 1$ to n may behave arbitrarily (which includes honest behaviour). This labelling of the peers captures the general case where some arbitrary $n - t$ peers may be dishonest, since the protocol is symmetric with respect to the labelling of the peers.

The model *BBProt* includes the Dolev-Yao adversary, and peers $t + 1$ to n considered to be under the control of the adversary. The setup is illustrated in Figure 6.

BBProt offers the same three external events as *BBSpec*, namely *post*, *ack* and *publish*. However it contains the peers explicitly, including peers controlled by the adversary, and so

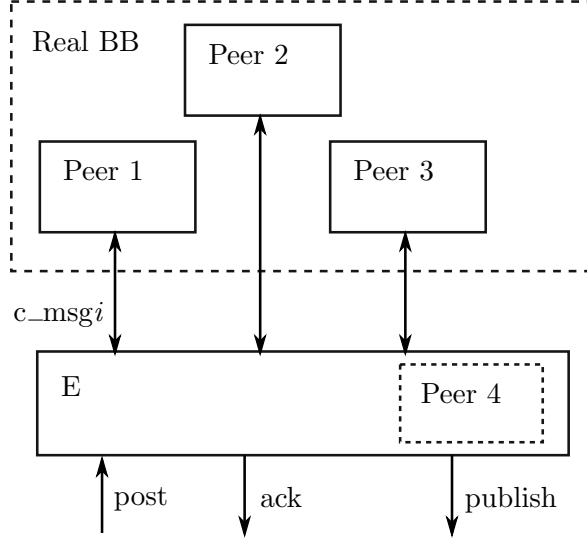


Figure 6: *BBProt*: Protocol model for analysis, with $t = 3$ and $n = 4$

the communication patterns with and between the peers will be quite different to those in the specification. Those communications are modelled by events $c_msg i$. The requirement for correctness will be that $BBSpec \preceq BBProt$ with respect to $\{post, ack, publish\}$.

The model *BBProt* includes the events that make up the various bulletin board protocols. Since these events can be performed whenever their guards are true, this means that interleavings of different protocols are naturally considered within this framework. Thus our approach to verification automatically allows for possible interference between the protocols, and a proof of correctness establishes that the protocols cannot interfere in an adverse way.

3.3 A useful lemma

The following lemma and corollary will be useful in the refinement proofs.

Lemma 3.1. *If $A \subseteq \{1, \dots, n\}$, $B \subseteq \{1, \dots, n\}$, $\#A \geq t$, $\#B \geq t$, and $t > 2n/3$, then there is some $j \leq t$ such that $j \in A$ and $j \in B$.*

Proof. We first prove that if $A \subseteq \{1, \dots, n\}$, $B \subseteq \{1, \dots, n\}$, $C \subseteq \{1, \dots, n\}$, $\#A \geq t$, $\#B \geq t$, $\#C \geq t$, and $t > 2n/3$, then $\#(A \cap B \cap C) \geq 1$.

We use the law $\#X + \#Y = \#(X \cup Y) + \#(X \cap Y)$. Observe that $A \cup B \subseteq \{1, \dots, n\}$ and so $\#(A \cup B) \leq n$. We obtain:

$$\begin{aligned} \#(A \cap B) &= \#A + \#B - \#(A \cup B) \\ &\geq t + t - n \end{aligned}$$

Then

$$\begin{aligned}
\#((A \cap B) \cap C) &= \#(A \cap B) + \#C - \#((A \cap B) \cup C) \\
&\geq (2t - n) + t - n \\
&= 3t - 2n
\end{aligned}$$

Now $t > 2n/3$, so $3t - 2n > 0$. Thus $\#((A \cap B) \cap C) \geq 1$ as required.

The result then follows immediately by setting $C = \{1, \dots, t\}$: then there is some $j \in A \cap B \cap C$, i.e. $j \leq t$ and $j \in A \cap B$.

Corollary 3.2. *If $A \subseteq \{1, \dots, t\}$, $B \subseteq \{1, \dots, t\}$, $\#A \geq 2t - n$, $\#B \geq 2t - n$, and $t > 2n/3$, then there is some $j \leq t$ such that $j \in A$ and $j \in B$.*

Proof. The corollary follows from Lemma 3.1 on $A \cup \{t + 1, \dots, n\}$ and $B \cup \{t + 1, \dots, n\}$.

4 One-shot Bulletin Board

To structure the analysis, we will consider the protocol in four stages:

1. Firstly we have a single posting phase and a single commit phase for publication of the items posted to the bulletin board.
2. We next introduce multiple commit phases for multiple updates of the published bulletin board.
3. Next we allow for the BB to reject some posts based on previous posts
4. Finally we optimise each commit phase to optimistic and fallback, using hash functions.

Our first model, introduced here, provides a one-shot bulletin board, which accepts posts for a period of time and then publishes its contents.

4.1 Specification

We model the specified behaviour in terms of the bulletin board communicating with its environment over a medium as illustrated in Figure 5. We now give definitions for the events within that framework. The given set $ITEM$ is the set of all items that can validly be posted to the bulletin board. In practice there will be some mechanism for recognising a valid post, such as a signature, but for the purposes of this paper we abstract such a mechanism and assume that only elements of $ITEM$ are posted. This corresponds to the expectation that posts not from $ITEM$ will be recognised and rejected by the bulletin board.

As described earlier, signatures are used to prevent the faking of receipts and the publishing of the bulletin board contents. The bulletin board uses (threshold) signature key SSK to sign receipts, and to sign the publication of the board. We define

$$\begin{aligned}
RECEIPT &= \{sig_{SSK}(x) \mid x \in ITEM\} \\
PUBLISH &= \{sig_{SSK}(B) \mid B \subseteq ITEM\}
\end{aligned}$$

```

machine BBSpec1
variables  $E_A, R, C$ 
invariant  $E_A \subseteq ITEM \cup RECEIPT \cup PUBLISH$ 
            $R \subseteq ITEM$ 
            $C \subseteq ITEM$ 

events
  init  $\hat{=}$   $E_A := \{\} \parallel R := \{\} \parallel C := \{\}$ ;
  post( $x$ )  $\hat{=}$  when  $x \in ITEM$  then  $E_A := E_A \cup \{x\}$  end;
   $r \leftarrow$  ack  $\hat{=}$   $r := (E_A \cap RECEIPT)$ ;
   $P \leftarrow$  publish  $\hat{=}$   $P := (E_A \cap PUBLISH)$ ;
  a_msg1  $\hat{=}$ 
    any  $x \in E_A \cap ITEM$ 
    then  $R := R \cup \{x\}$ 
    end;
  a_msg2  $\hat{=}$ 
    any  $x$ 
    where  $x \in R \wedge (sig_{SSK}(B) \in E_A \Rightarrow x \in B)$ 
    then  $E_A := E_A \cup \{sig_{SSK}(x)\} \parallel C := C \cup \{x\}$ 
    end;
  a_msg3  $\hat{=}$ 
    any  $Y$ 
    where  $C \subseteq Y \subseteq R \wedge E_A \cap PUBLISH = \{\}$ 
    then  $E_A := E_A \cup \{sig_{SSK}(Y)\}$ 
    end
end

```

Figure 7: Bulletin Board Model incorporating the environment

Allowing for untrusted peers requires us to include some nondeterminism within the specification of the bulletin board, to reflect (bb.1) and (bb.2) above. In particular, dishonest peers and the untrusted medium can prevent receipts from being issued for some received posts, so the specification must allow for this possibility. The model of the bulletin board thus uses two databases: R consisting of received posts, and C consisting of confirmed posts—those which have been acknowledged with receipts.

When the board B is published, anything published must be in R in accordance with (bb.1); and all confirmed posts C must be published in accordance with (bb.2). Thus we require $C \subseteq B \subseteq R$. In other words, items that have been submitted to the bulletin board but not confirmed might or might not appear in B . We retain a level of uncertainty over what is published, because this level of uncertainty is present in the implementation when some of the bulletin board peers are untrusted. Furthermore, in the implementation the adversary can orchestrate further posts and receipts following publication of the bulletin board, so our specification must reflect this: additional posts can be accepted. Requirement (bb.2) states that given both a published bulletin board and a receipted item, that item must be on the bulletin board. To remain consistent with this requirement, any receipts issued in a_msg2 after bulletin board publication must be on any published bulletin board. Observe that if every posting has a receipt, then the bulletin board will contain all posted items ($C = R = B$).

Observe that a_msg3 allows no more than one bulletin board to be published, meeting requirement (bb.4): once published, the bulletin board is fixed.

The resulting Event-B model is given in Figure 7. This is the specification that we will show our design meets. The model includes a bulletin board and its environment. As well as the state of the bulletin board, we include the state E_A of the environment, containing the communications that it is managing, because we will want to consider the bulletin board protocol design in a model including the Dolev-Yao adversary, which provides an asynchronous communication medium.

4.2 Implementation: a Robust Bulletin Board Design

The aim of the implementation is that if a threshold of peers behave according to the protocol, then the implementation will behave as the bulletin board of Section 4 above with receipts and publication commitment. This allows for a minority of peers to fail, or to behave maliciously, without impacting on the overall behaviour of the bulletin board. In fact as we shall see, as long as a post x is handled by some threshold of peers then a receipt can be provided, and x will appear on the public web bulletin board. Different posts can be handled by different threshold sets, allowing for individual peers to drop out temporarily (e.g. from a temporary loss of communication).

There are n peers, numbered 1 to n . Each peer j has its own signing key sk_j . There is also a threshold signing key SSK , and each peer j has a share of it: ssk_j . Any t out of n partial signatures SSK on a value m can be combined to the corresponding signature on m : $sig_{SSK}(m)$. The condition on the threshold t is that $t > 2n/3$.

The design is a slight simplification of that given in [Sur13]. Each peer j maintains its local database D_j , which initially contains no entries. It also has a boolean variable pub_j which is initially *false*.




Figure 8: Posting Protocol

The peers run two protocols. The first is for accepting posts and providing acknowledgements, and the second is for publishing the bulletin board.

Post and Acknowledge Protocol

This protocol is illustrated in Figure 8, and is a simplification of the full protocol of Figure 1. It consists of three rounds, as follows:

1. $User \rightarrow P_i : x$ (for each $i \in I$)
2. $P_i \rightarrow P_j : sig_{sk_i}(x)$ (for each $i, j \in I, j \neq i$)
each P_i waits for at least a threshold number of signatures
3. $P_i \rightarrow User : sig_{ssk_i}(x)$ (for each $i \in I$)

Publish Protocol

This protocol is illustrated in Figure 9. It is a combination of the pair of protocols given in Figures 2 and 3, with Round 1 as the fallback protocol and then Round 2 as the optimistic




Figure 9: Publishing Protocol

protocol, and with the signature on the bulletin board directly. When the time comes to publish, then the Post and Acknowledge protocol stops and is no longer executed, and the peer begins the commit protocol which is used for the peers to obtain agreement on the bulletin board to publish, as follows:

1. $P_i \rightarrow P_j$: D_i (for each $i, j \in I, j \neq i$)
each P_i adds any missing information received from others to its own database
2. $P_i \rightarrow WBB$: $sig_{ssk_i}(t(D_i))$ (for each $i \in I$)

4.3 Event-B Modelling and analysis

The threat model built into the model of the protocol incorporates our robustness considerations, in particular that the correctness of the bulletin board is not dependent on the correct behaviour of any individual component, as long as a threshold behave correctly. It allows for the case where peers behave honestly but occasionally are down (either through connection loss, or through temporary server loss): this is modelled simply by the absence of messages between DY and the peer, and allows for peers to miss some posts. The model also includes the case where peers $t + 1$ to n can lose or otherwise alter their databases of received posts. However, the honest peers 1 to t do not lose their databases: for correctness we require that a threshold of peers do not lose their data.

The set *MESSAGE* of all possible messages m in the model is given as follows:

$$m ::= k \mid i \mid sig_k(m) \mid \{m_1, \dots, m_n\}$$

where $k \in KEY$ and $i \in ITEM$. Observe that a message can itself consist of a set of messages, and thus *MESSAGE* covers all rounds of the protocol. In particular *RECEIPT* \subseteq *MESSAGE* and *PUBLISH* \subseteq *MESSAGE*.

Two further definitions will be useful when expressing the model:

$$\begin{aligned} SIG1 &= \{sig_{sk_k}(x) \mid 1 \leq k \leq n \wedge x \in ITEM\} \\ t(D) &= \{x \mid \#\{k \mid sig_{sk_k}(x) \in D\} \geq t\} \end{aligned}$$

The set *SIG1* is the set of items signed by any of the peers. Given a set *D* of signed items, the set $t(D)$ is those items for which *D* contains a threshold number of different signatures. If *D* is used to track the signed items received by a peer, then $t(D)$ is those items for which it has received a threshold number.

We will now define the model. It is declared as follows:

```

machine BBProt1
refines BBSpec1
variables E, Ij, Dj, pubj, comj ( $1 \leq j \leq t$ )

```

Its invariant is given as follows:

```

invariant
/* Types */
   $E \subseteq MESSAGE$ 
   $I_j \subseteq ITEM$ 
   $D_j \subseteq SIG1$ 
   $pub_j \in BOOL$ 
   $com_j \in BOOL$ 
/* Key invariant properties */
   $k \leq t \wedge k \in c[x] \wedge k \in s[B] \Rightarrow x \in B$  (1)
   $sig_{ssk_j}(x) \in E \Rightarrow \#d_j[x] \geq t$  (2)
   $sig_{SSK}(x) \in E \Rightarrow \#c[x] \geq t$  (3)
   $sig_{ssk_j}(B) \in E \Rightarrow B \subseteq t(D_j)$  (4)
   $sig_{SSK}(B) \in E \Rightarrow \#s[B] \geq t$  (5)
   $D_j \subseteq E$  (6)
   $k \leq t \wedge k \in s[B] \Rightarrow com_k = true$  (7)
   $k \leq t \wedge k \in s[B_1] \wedge B_1 \neq B_2 \Rightarrow k \notin s[B_2]$  (8)

/* adversary bound invariant — see (9)–(11) below
/* Linking invariant — see (12)–(14) of Section 4.4 */
where:
   $d_j[x] = \{k \mid sig_{sk_k}(x) \in D_j\}$  shares of part sigs on  $x$  received by Peer  $j$ 
   $c[x] = \{k \mid sig_{ssk_k}(x) \in E\}$  peers which have (part)signed the receipt on  $x$ 
   $s[B] = \{k \mid sig_{ssk_k}(B) \in E\}$  peers which have part-signed bulletin board  $B$ 

```

Each event introduced below preserves the invariant: **INV** is established for each event.

The initialisation and external events are given as follows:

```

events
  init  $\hat{=}$   $E := \{sk_k \mid k > t\} \cup \{ssk_k \mid k > t\} \parallel$ 
     $\parallel_j (I_j := \{\} \parallel D_j := \{\} \parallel pub_j := false \parallel com_j := false);$ 

  post( $x$ )  $\hat{=}$  when  $x \in ITEM$  then  $E := E \cup \{x\}$  end;

   $r \leftarrow ack \hat{=}$   $r \in (E_A \cap RECEIPT);$ 

   $P \leftarrow publish \hat{=}$   $P \in (E_A \cap PUBLISH);$ 

```

The initial state of each *Peer* j has $I_j = \emptyset$, $D_j = \emptyset$, $pub_j = false$ and $com_j = false$.

Post and Acknowledge

Recall that Round 2 of the posting protocol involves each peer sending a message to all the other peers. This is split into two events: c_msg2a for the sending of the message (to be held

by the Dolev-Yao environment), and c_msg2b for peers receiving the message. We model the sending of the message to all peers by sending the message to the environment, and then allowing all the other peers to receive it.

c_msg1: $DY \rightarrow Peer\ j : x$

If $\neg pub_j$, then *Peer j* inputs the post x . *Peer j* adds x to its local database of received items I_j .

```

c_msg1_j  $\hat{=}$ 
  any  $x \in E \cap ITEM \wedge \neg pub_j$ 
  then  $I_j := I_j \cup \{x\}$ 
  end;

```

c_msg2a: $Peer\ j \rightarrow DY : sig_{sk_j}(x)$ [*Peer j* to all other peers.]

If $\neg pub_j$, and I_j contains x then *Peer j* creates $sig_{sk_j}(x)$ with its signature key, outputs it to DY intended for the other peers, and adds it to its local database D_j .

```

c_msg2a_j  $\hat{=}$ 
  any  $x$ 
  where  $x \in I_j \wedge \neg pub_j$ 
  then  $E := E \cup \{sig_{sk_j}(x)\} \parallel D_j := D_j \cup \{sig_{sk_j}(x)\}$ 
  end;

```

c_msg2b: $DY \rightarrow Peer\ j : sig_{sk_k}(x)$ [*Peer j* inputting from other peers.]

If $\neg pub_j$, and *Peer j* inputs $sig_{sk_k}(x)$, then *Peer j* adds $sig_{sk_k}(x)$ to D_j .

```

c_msg2b_j  $\hat{=}$ 
  any  $x, k$ 
  where  $sig_{sk_k}(x) \in E \wedge \neg pub_j$ 
  then  $D_j := D_j \cup \{sig_{sk_k}(x)\}$ 
  end;

```

c_msg3: $Peer\ j \rightarrow DY : sig_{ssk_j}(x)$

If $\neg pub_j$, and D_j contains t different signatures on x , then *Peer j* outputs a signature share $sig_{ssk_j}(x)$.

```

c_msg3_j  $\hat{=}$ 
  any  $x$ 
  where  $x \in t(D_j)$ 
  then  $E := E \cup \{sig_{ssk_j}(x)\}$ 
  end;

```

When DY has a threshold number of signature shares $sig_{ssk_j}(x)$ on x , DY can combine them to form the receipt $sig_{SSK}(x)$, and add this to E . (See event c_dy2 below.)

Commit and Publish

The publish protocol starts by setting pub_j to true:

c_msg4: *Peer j : commit_j*

```

c_msg4j  $\hat{=}$ 
  when  $\neg pub_j$ 
  then  $pub_j := true$ 
  end;

```

Round 1 of the protocol, each peer sending a signature share to all the others, is modelled by two events: sending, and receiving.

c_msg5a: *Peer j \rightarrow DY : B_j* [*Peer j to all other peers.*]
 If pub_j then *Peer j* outputs B_j , its local database of signed items, intended for the other peers. This communicates its local database to the other peers.

```

c_msg5aj  $\hat{=}$ 
  when  $pub_j$ 
  then  $E := E \cup \{D_j\}$ 
  end;

```

c_msg5b: *DY \rightarrow Peer j : D_k*
 If pub_j then *Peer j* inputs D_k , k 's local database D_k . This is added to D_j : any signed posts $sig_{sk_k}(x)$ in D_k that are not already in D_j are added to D_j .

```

c_msg5bj  $\hat{=}$ 
  any  $D$ 
  where  $D \in E \wedge D \subseteq SIG1 \wedge pub_j$ 
  then  $D_j := D_j \cup D$ 
  end;

```

c_msg6: *Peer j \rightarrow DY : sig_{ssk_j}(t(D_j))*
 If pub_j then *Peer j* can send out a signature share on its current version of the bulletin board: those items for which it holds a threshold of signatures, $t(D_j)$.

```

c_msg6j  $\hat{=}$ 
  when  $pub_j \wedge \neg com_j$ 
  then  $E := E \cup \{sig_{ssk_j}(t(D_j))\} \parallel com_j := true$ 
  end;

```

The Dolev-Yao environment

The Dolev-Yao environment is modelled through the use of the set E to retain all messages that are sent and received by protocol parties. The adversary is also able to generate new messages to introduce into protocol executions. In particular, he can sign any message with any key that he possesses; he can combine shares of a signature into a threshold signature; he can extract the message from a signature; and he can add and remove messages from a set of messages. These capabilities are captured in the following derivation rules, which show how a new message can be generated from a set of messages.

$$\begin{aligned}
& \{k, m\} \vdash \text{sig}_k(m) \\
\#S \geq t \Rightarrow \quad & \{\text{sig}_{ssk_k}(m) \mid k \in S\} \vdash \text{sig}_{SSK}(m) \\
& \{\text{sig}_k(m)\} \vdash m \\
& \{m, B\} \vdash B \cup \{m\} \\
m \in B \Rightarrow \quad & \{B\} \vdash m
\end{aligned}$$

We model adversary behaviour by including an event for each rule, allowing the adversary to introduce new events to the set E .

```

c_dy1  $\hat{=}$       /* signing */
  any  $m, s$ 
  where  $m \in E \wedge s \in E$ 
  then  $E := E \cup \{\text{sig}_s(m)\}$ 
  end;

c_dy2  $\hat{=}$       /* threshold signature on  $m$  */
  any  $S, m$ 
  where  $\#S \geq t \wedge \{\text{sig}_{ssk_k}(m) \mid k \in S\} \subseteq E$ 
  then  $E := E \cup \{\text{sig}_{SSK}(m)\}$ 
  end;

c_dy3  $\hat{=}$       /* extracting  $m$  from signature */
  any  $m, s$ 
  where  $\text{sig}_s(m) \in E$ 
  then  $E := E \cup \{m\}$ 
  end;

c_dy4  $\hat{=}$       /* adding  $m$  to  $B$  */
  any  $m, B$ 
  where  $m \in E \wedge B \in E$ 
  then  $E := E \cup \{B \cup \{m\}\}$ 
  end;

c_dy5  $\hat{=}$       /* extracting  $m$  from  $B$  */
  any  $m, B$ 
  where  $B \in E \wedge m \in B$ 
  then  $E := E \cup \{m\}$ 
  end;

```

Some additional clauses are necessary to introduce into the invariant as below, to capture the limits of what the adversary can introduce. These are necessary for the refinement proof.

/ adversary bound invariant */*

$$E \cap (\{sk_k \mid k \leq t\} \cup \{ssk_k \mid k \leq t\}) = \emptyset \quad (9)$$

$$\bigcup_{e \in E} items(e) \subseteq E \quad (10)$$

$$\bigcup_{e \in E} sigs(e) \subseteq E \quad (11)$$

where

$$\begin{aligned} items(x) &= \{x\} \\ items(sig_s(x)) &= items(x) \\ items(B) &= \bigcup_{b \in B} items(b) \\ sigs(x) &= \{\} \\ sigs(sig_s(x)) &= \{sig_s(x)\} \cup sigs(x) \\ sigs(B) &= \bigcup_{b \in B} sigs(b) \end{aligned}$$

4.4 Simulation

We aim to establish that the concrete system *BBProt1* refines the abstract system *BBSpec1* with respect to the external events $\{post, ack, publish\}$.

To establish refinement we show that any concrete move can be matched by an abstract move, or (for events other than *post*, *ack* and *publish*) matched by *skip*. To do this we need to identify the *linking invariant*, the relationship between the abstract and concrete states, and show that any concrete move from a concrete state is matched for any corresponding abstract state by some abstract move or *skip*.

Linking invariant

We thus have to identify when, in the concrete system, abstract events are considered to have occurred.

- abstract *a_msg1* occurs when the bulletin board receives x . In the concrete model this corresponds to t peers having received x and signed it. Since there can be up to $n - t$ dishonest peers, this means $t - (n - t) = 2t - n$ honest peers having signed x .

- abstract a_msg2 occurs when the bulletin board issues a signature on x . This corresponds to the combining of t returns of signature shares $sig_{ssk_j}(x)$.
- abstract a_msg3 occurs when a signed database $sig_{SSK}(t(D))$ is produced. This corresponds to the combining of t returns of signature shares $sig_{ssk_j}(t(D))$.

The abstract state s_A is the pair of databases R and C , and medium E_A .

The concrete state s_C is the set of databases I_j , D_j and pub_j for the peers, E for the Dolev-Yao environment.

The linking invariant is given by the following predicate $J(s_A, s_C)$:

/* linking invariant */

$$R = \{x \in ITEM \mid \#\{k \mid 1 \leq k \leq t \wedge sig_{ssk_k}(x) \in E\} \geq 2t - n\} \quad (12)$$

$$C = \{x \in ITEM \mid sig_{SSK}(x) \in E\} \quad (13)$$

$$E_A = E \cap (ITEM \cup RECEIPT \cup PUBLISH) \quad (14)$$

R is the set of items for which the adversary (and possibly other peers) can provide a threshold of $sig_{ssk_j}(x)$, and so can include x on the published bulletin board. If at least $2t - n$ honest peers have signed x , then it is within the adversary's control to produce a further $n - t$ signatures, giving a threshold of signatures on x . C is the set of items for which the adversary has a receipt—evidence that sufficiently many peers have a threshold of $sig_{ssk_j}(x)$ to ensure that it will appear on the published bulletin board.

We are now in a position to present the main result: that the concrete model behaves according to the abstract model.

Lemma 4.1. *$BBSpec1 \preceq BBProt1$ with respect to $\{post, ack, publish\}$.*

Proof Consider each event of $BBProt1$ in turn. It is necessary to prove **GRD_INV_REF**₃ in each case. In most cases the event is matched by *skip* and we establish **INV_REF**₂, which is stronger.

Case *post*. Matched by *post* of $BBSpec1$: the update to E is matched by the update to E_A , preserving the linking invariant.

Case *ack*. Matched by *ack* of $BBSpec1$: the concrete output of receipt r is matched by the abstract output of receipt r , since if $r \in E$ then $r \in E_A$ by the linking invariant.

Case *publish*. Matched by *publish* of $BBSpec1$: the concrete output of M from E is matched by the abstract output of M , since if $M \in E$ then $M \in E_A$ by the linking invariant.

Case *c_msg1*. Matched by *skip*

Case *c_msg2a*. If $\#\{k \mid 1 \leq k \leq t \wedge sig_{ssk_k}(x) \in E\} = 2t - n - 1$ and $\#\{k \mid 1 \leq k \leq t \wedge sig_{ssk_k}(x) \in E'\} = 2t - n$ then this event is matched by $m_A = a_msg1$ with x . Otherwise matched by *skip*.

Case *c_msg2b*. Matched by *skip*.

Case c_msg3 . Matched by *skip*.

Case c_msg4 . Matched by *skip*.

Case c_msg5_a . Matched by *skip*.

Case c_msg5_b . Matched by *skip*.

Case c_msg6 . Matched by *skip*.

Case c_dy1 . Matched by *skip*. In particular, R remains unchanged due to invariant (9).

Case c_dy2 . For variable E , we use E to refer to its value before the occurrence of this event, and E' for its value after its occurrence.

If $x \in ITEM$ and $sig_{SSK}(x) \notin E$ and $sig_{SSK}(x) \in E'$, then this is matched by a_msg2 . We must show that (1) $x \in R$ and (2) $sig_{SSK}(B) \in E_A \Rightarrow x \in B$.

1. $x \in R$: We have that $\#c[x] \geq t$. Hence there is some $k \leq t$ with $k \in c[x]$, so by invariant (2) it follows that $\#d_k[x] \geq t$. By invariant (6) it follows that $\#\{k \mid sig_{sk_k}(x) \in E\} \geq t$, and hence that $\#(\{k \mid sig_{sk_k}(x) \in E\} - \{t+1 \dots n\}) \geq t - (n-t) = 2t-n$. Hence $x \in R$ as required.
2. $sig_{SSK}(B) \in E_A \Rightarrow x \in B$: Assume $sig_{SSK}(B) \in E_A$. Then $\#s[B] \geq t$. Also we have $\#c(x) \geq t$, so by Lemma 3.1 there is some $k \leq t$ with $k \in c[x]$ and $k \in s[B]$. Hence from (1) it follows that $x \in B$ as required.

If $B_C \subseteq ITEM$ and $sig_{SSK}(B_C) \notin E$ and $sig_{SSK}(B_C) \in E'$, then this is matched by a_msg3 , with $B = B_C$. We must show that (1) $E_A \cap PUBLISH = \{\}$, (2) $C \subseteq B_C$ and (3) $B_C \subseteq R$.

1. $E_A \cap PUBLISH = \{\}$: we establish this by contradiction. If $sig_{SSK}(B) \in E_A$ for some $B \neq B_C$, then $\#s[B] \geq t$ by (5). Also we have $s[B_C] \geq t$ by the guard of c_dy2 . Hence from Lemma 3.1 there is some $k \leq t$ with $k \in s[B]$ and $k \in s[B_C]$, contradicting (8).
2. $C \subseteq B_C$: consider some $x \in C$. Then $sig_{SSK}(x) \in E$, so $\#c[x] \geq t$ by invariant (3). Further, $\#s[B_C] \geq t$ by invariant (5). Hence from Lemma 3.1 there is some $k \leq t$ with $k \in c[x]$ and $k \in s[B_C]$. Hence by invariant (1), $x \in B_C$, as required.
3. $B_C \subseteq R$: We have from invariant (3) that $sig_{ssk_j}(B_C) \in E$ for some $j \leq t$. Now consider $x \in B_C$. Then $x \in t(D_j)$ by invariant (4). Hence $x \in t(E)$ by invariant (6), and so $\#\{k \mid 1 \leq k \leq n \wedge sk_k(x) \in E\} \geq t$ from the definition of $t(E)$, and hence $\#\{k \mid 1 \leq k \leq t \wedge sk_k(x) \in E\} \geq 2t-n$. Thus $x \in R$ as required.

Otherwise c_dy2 is matched by *skip*.

Case c_dy3 . Matched by *skip*, since $E_A = E \cap (ITEM \cup RECEIPT \cup PUBLISH)$ does not change, by invariants (10) and (11).

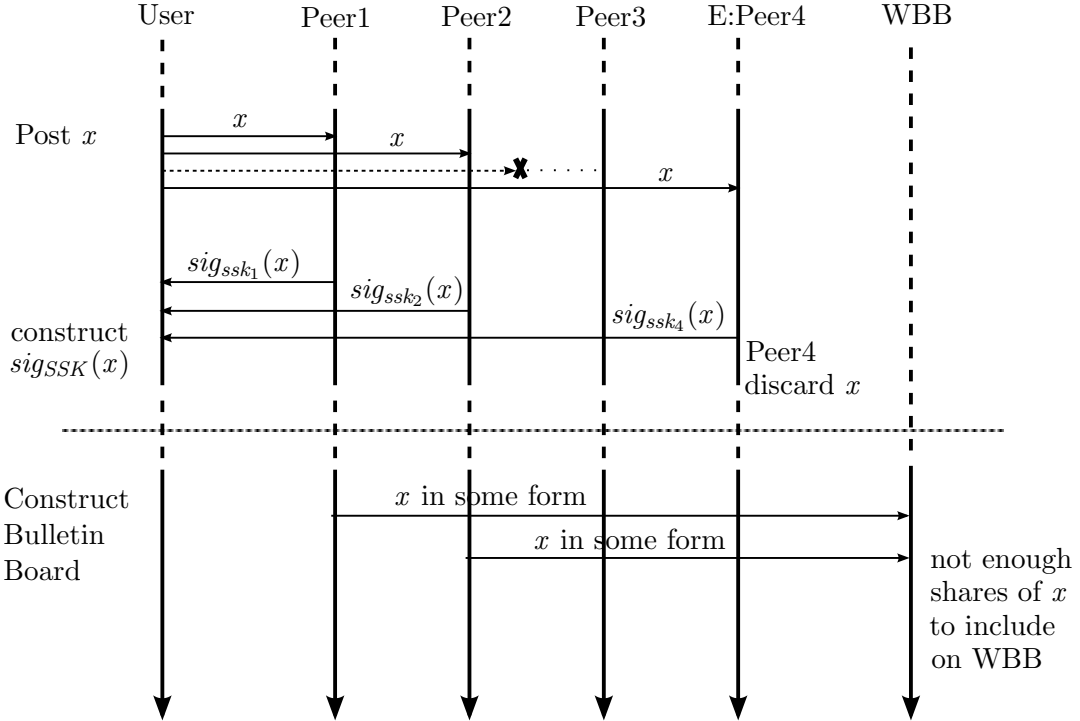
Case c_dy4 . Matched by *skip*.

Case c_dy5 . Matched by *skip*, since any items, receipts or publish messages in B are already in E by invariants (10) and (11).

This concludes the proof that $BBSpec1 \preceq BBProt1$ with respect to $\{post, ack, publish\}$, establishing the correctness of the bulletin board protocol $BBProt1$ against the specification $BBSpec1$.

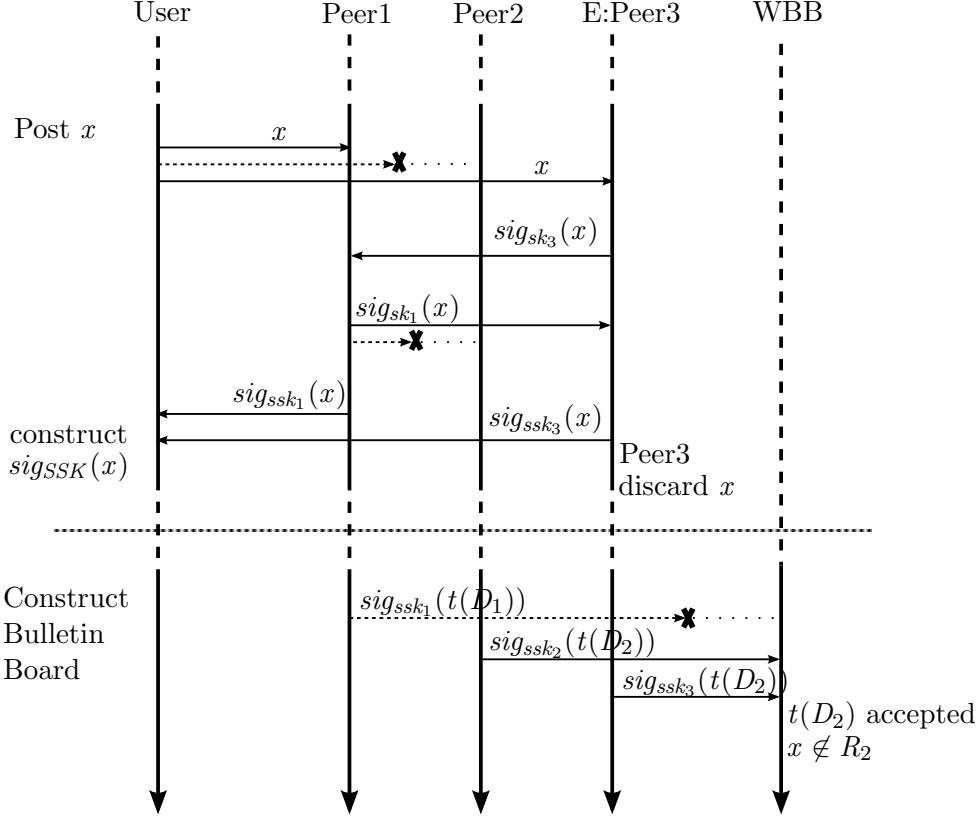
4.5 Example attacks on weaker versions

Example 4.1. In order to see the necessity for the round of signed messages in the posting and acknowledgement protocol (message 2), we consider what can occur if this round is not included. In particular, if peers simply receive posts and respond with their signature share towards the receipt, then an adversary can organise for a receipt to be provided for an item not on the bulletin board, as follows:



Here Peer 3 is cut out of the posting and acknowledgement protocol on a submission x , but a threshold of peers provide a signature share and so a receipt is provided. However, peer 4 is dishonest, and so discards x before publication of the bulletin board. Hence there are only two shares of x recorded, insufficient to warrant inclusion on the published bulletin board.

Example 4.2. This example illustrates the necessity for the threshold to be greater than $2n/3$. If the threshold is $2n/3$ or less, then an adversary can arrange for a receipt to be issued on an item not included on the bulletin board, as follows:



In this attack a receipt can be provided on a post x although it does not appear on the board. Peer 2 is excluded from the posting and acknowledgement of x , however participation from Peers 1 and 3 is sufficient to provide a receipt. Peer 3 (which is dishonest) then discards x . When the bulletin board is published, Peer 1 is excluded from the publication protocol, but Peers 2 and 3 agree on a bulletin board not including x , and so that is published. The attack works because there is no honest peer that has participated in both the acknowledgement of x and its posting on the bulletin board. The attack cannot happen if the threshold is strictly greater than $2n/3$, because in that case there must be some honest peer contributing to both the receipt on x and the agreed bulletin board, which is enough to ensure that x is included on the bulletin board.

5 Multiple Bulletin Board Rounds

We extend to the case where multiple bulletin boards can be published. We consider a period p to consist of a number of posts followed by publication of the associated bulletin board for that period. Thus different bulletin boards can be published for different periods, and we require that every period's bulletin board will behave according to the bulletin board specification given in *BBSpec1*.

5.1 Specification

The specification of multiple bulletin boards is of a collection of boards that each behave according to specification *BBSpec1*. This is captured as an indexed collection of bulletin boards within a single specification *BBSpec2*. Receipts will be issued with the index of the bulletin board the item has been posted to, and a bulletin board will be published with its index. We define

$$\begin{aligned} RECEPT2 &= \{sig_{SSK}(p, x) \mid x \in ITEM \wedge p \in \mathbb{N}\} \\ PUBLISH2 &= \{sig_{SSK}(p, B) \mid B \subseteq ITEM \wedge p \in \mathbb{N}\} \\ PUBLISH2_p &= \{sig_{SSK}(p, B) \mid B \subseteq ITEM\} \end{aligned}$$

```

machine BBSpec2
variables  $E_A, R_p, C_p \quad (p \in \mathbb{N})$ 
invariant  $E_A \subseteq ITEM \cup RECEPT2 \cup PUBLISH2$ 
            $R_p \subseteq ITEM \quad (p \in \mathbb{N})$ 
            $C_p \subseteq ITEM \quad (p \in \mathbb{N})$ 
events
  init  $\hat{=}$   $E_A := \{\} \parallel \parallel_{p \in \mathbb{N}} (R_p := \{\} \parallel C_p := \{\})$ 
  post( $x$ )  $\hat{=}$  when  $x \in ITEM$  then  $E_A := E_A \cup \{x\}$  end;
   $r \leftarrow \text{ack} \hat{=}$   $r : \in (E_A \cap RECEPT2)$ ;
   $P \leftarrow \text{publish} \hat{=}$   $P : \in (E_A \cap PUBLISH2)$ ;
  a_msg1  $\hat{=}$ 
    any  $x \in E_A \cap ITEM \wedge p \in \mathbb{N}$ 
    then  $R_p := R_p \cup \{x\}$ 
    end;
  a_msg2  $\hat{=}$ 
    any  $x, p$ 
    where  $x \in R_p \wedge (sig_{SSK}(p, B) \in E_A \Rightarrow x \in B)$ 
    then  $E_A := E_A \cup \{sig_{SSK}(p, x)\} \parallel C_p := C_p \cup \{x\}$ 
    end;
  a_msg3  $\hat{=}$ 
    any  $Y, p$ 
    where  $C_p \subseteq Y \subseteq R_p \wedge E_A \cap PUBLISH2_p = \{\}$ 
    then  $E_A := E_A \cup \{sig_{SSK}(p, Y)\}$ 
    end
end

```

5.2 Implementation

In the implementation, each peer maintains a counter p_j which it uses to track the period it is currently accepting posts for. The counter will be incremented when it has finished accepting

posts for one period and begins accepting posts for the next. It also maintains a separate state space for each period. For example, where *BBProt1* used R_j for j 's record of what it had received, *BBProt2* will use $R_{j,p}$ for j 's record of what it received in period p , and so will have a separate set for each period.

The resulting model *BBProt2* is given in the various clauses below. The model is shown in the events within the description. The key to the refinement proof is that the interleaving of the events across the different periods do not interfere, even though peers can progress their periods independently and can be involved in publication of one bulletin board while receiving items for another.

Declaration and Invariant

$$SIG1_p = \{sig_{SSK}(p, x) \mid x \in ITEM\}$$

```

machine BBProt2
refines BBSpec2
variables  $E, I_{j,p}, D_{j,p}, p_j \ (1 \leq j \leq t)$ 

invariant
/* Types */
 $E \subseteq MESSAGE$ 
 $I_{j,p} \subseteq ITEM$ 
 $D_{j,p} \subseteq \{sig_{sk_k}(p, x) \mid x \in ITEM\}$ 
 $p_j \in \mathbb{N}$ 
/* Key invariant properties */
 $k \leq t \wedge k \in c[p, x] \wedge k \in s[p, B] \Rightarrow x \in B$ 
 $sig_{ssk_j}(p, x) \in E \Rightarrow \#d_j[p, x] \geq t$ 
 $sig_{SSK}(p, x) \in E \Rightarrow \#c[p, x] \geq t$ 
 $sig_{ssk_j}(p, B) \in E \Rightarrow B \subseteq t(D_{j,p})$ 
 $sig_{SSK}(p, B) \in E \Rightarrow \#s[p, B] \geq t$ 
 $D_{j,p} \subseteq E$ 
 $k \leq t \wedge k \in s[p, B] \Rightarrow c_k > p$ 
 $k \leq t \wedge k \in s[p, B_1] \wedge B_1 \neq B_2 \Rightarrow k \notin s[p, B_2]$ 
/* linking invariant */
 $R_p = \{x \in ITEM \mid \#\{k \mid 1 \leq k \leq t \wedge sig_{sk_k}(p, x) \in E\} \geq 2t - n\}$ 
 $C_p = \{x \in ITEM \mid sig_{SSK}(p, x) \in E\}$ 
 $E_A = E \cap (ITEM \cup RECEIPT2 \cup PUBLISH2)$ 

where:
 $d_j[p, x] = \{k \mid sig_{sk_k}(p, x) \in D_j\}$       shares of part sigs on  $x$  received by Peer  $j$ 
 $c[p, x] = \{k \mid sig_{ssk_k}(p, x) \in E\}$       peers which have (part)signed the receipt on  $x$ 
 $s[p, B] = \{k \mid sig_{ssk_k}(p, B) \in E\}$       peers which have part-signed bulletin board  $B$ 

```

External events

External events look very similar in *BBProt2*.

```

events
  init  $\hat{=}$   $E := \{sk_k \mid k > t\} \cup \{ssk_k \mid k > t\} \parallel$ 
     $\parallel_{j,n} (I_{j,n} := \{\} \parallel D_{j,n} := \{\} \parallel p_j := 0 \parallel c_j := 0);$ 
  post( $x$ )  $\hat{=}$  when  $x \in ITEM$  then  $E := E \cup \{x\}$  end;
   $r \leftarrow \text{ack} \hat{=}$   $r := (E \cap RECEIPT2);$ 
   $P \leftarrow \text{publish} \hat{=}$   $P := E \cap PUBLISH2;$ 

```

Posting and acknowledgement protocol

Posting and acknowledgement is similar. The new aspect is the introduction of the period p_j , and $Peer_j$ may only accept and acknowledge items, and issue its share of the receipt, for items in its current period.

```

c_msg1_j( $x$ )  $\hat{=}$  /* receive item  $x$  */
  when  $x \in E \cap ITEM$ 
  then  $I_{j,p_j} := I_{j,p_j} \cup \{x\}$ 
  end;

c_msg2a_j  $\hat{=}$  /* send signature share on  $x$  */
  any  $x$ 
  where  $x \in I_{j,p_j}$ 
  then  $E := E \cup \{sig_{sk_j}(p_j, x)\} \parallel D_{j,p_j} := D_{j,p_j} \cup \{sig_{sk_j}(p_j, x)\}$ 
  end;

c_msg2b_j  $\hat{=}$  /* receive signature share on  $x$  */
  any  $x, k$ 
  where  $sig_{sk_k}(p_j, x) \in E$ 
  then  $D_{j,p_j} := D_{j,p_j} \cup \{sig_{sk_k}(p_j, x)\}$ 
  end;

c_msg3_j  $\hat{=}$  /* send signature share on receipt of  $x$  */
  any  $x$ 
  where  $x \in t(D_{j,p_j})$ 
  then  $E := E \cup \{sig_{ssk_j}(p_j, x)\}$ 
  end;

```

Commit and publish protocol

The commit protocol for the bulletin board of period p_j is started by incrementing p_j . Thus no further posts will be accepted for that bulletin board, and the events in the commit and publish protocol are then enabled. They match the events from *BBProt1*.

```

c_msg4j  $\hat{=}$       /* start commit protocol */
begin
   $p_j := p_j + 1$ 
end;

c_msg5aj  $\hat{=}$       /* send database */
  any  $p < p_j$ 
  then  $E := E \cup \{D_{j,p}\}$ 
end;

c_msg5bj  $\hat{=}$       /* receive  $k$ 's database, update  $D_{j,p}$  if necessary */
  any  $D, p$ 
  where  $D \in E \wedge D \subseteq \text{SIG1}_p \wedge p < p_j$ 
  then  $D_{j,p} := D_{j,p} \cup D$ 
end;

c_msg6j  $\hat{=}$       /* publish signature share on  $t(D_{j,p})$  */
  when  $c_j < p_j$ 
  then  $E := E \cup \{\text{sig}_{ssk_j}(c_j, t(D_{j,c_j}))\} \parallel c_j := c_j + 1$ 
end;

```

Dolev-Yao environment

The adversary has the same moves as before, with two new ones, combining and separating pairs. This arises from the introduction of pairing in this model, to allow the period along with the message to be signed.

<pre> c_dy1 $\hat{=}$ /* signature share on m */ any m, s where $m \in E \wedge s \in E$ then $E := E \cup \{sig_s(m)\}$ end; c_dy2 $\hat{=}$ /* threshold signature on m */ any S, m where $\#S \geq t \wedge \{sig_{ssk_k}(m) \mid k \in S\} \subseteq E$ then $E := E \cup \{sig_{SSK}(m)\}$ end; c_dy3 $\hat{=}$ /* extracting m from signature */ any m, s where $sig_s(m) \in E$ then $E := E \cup \{m\}$ end; </pre>	<pre> c_dy4 $\hat{=}$ /* adding m to B */ any m, B where $m \in E \wedge B \in E$ then $E := E \cup \{B \cup \{m\}\}$ end; c_dy5 $\hat{=}$ /* extracting m from B */ any m, B where $B \in E \wedge m \in B$ then $E := E \cup \{m\}$ end; c_dy6 $\hat{=}$ /* pairing */ any m, p where $m \in E \wedge p \in \mathbb{N}$ then $E := E \cup \{(p, m)\}$ end; c_dy7 $\hat{=}$ /* splitting */ any m, p where $(p, m) \in E$ then $E := E \cup \{p, m\}$ end </pre>
--	--

5.3 Simulation

Establishing the simulation relation follows the structure of the proof that *BBProt1* refines *BBSpec1*. *BBProt2* essentially consists of an indexed collection of *BBProt1* bulletin boards. Each peer j maintains a counter p_j indicating its current bulletin board. The bulletin board indexed by p has $p < p_j$ in place of pub_j : Peer j enters the publication phase for bulletin board p once the counter p_j has progressed beyond p . It also has $p < c_j$ in place of com_j : Peer j has committed to its share once the counter c_j has progressed beyond p .

Thus we obtain:

Lemma 5.1. *BBSpec2* \preceq *BBProt2* with respect to $\{post, ack, publish\}$

Proof (sketch)

We need to prove that if $J(s_A, s_C)$, and $s_C \xrightarrow{m_C} s'_C$ then either $J(s_A, s'_C)$ (m_C is matched by *skip*), or $\exists m_A, s'_A$ such that $s_A \xrightarrow{m_A} s'_A$ and $J(s'_A, s'_C)$ (m_C is matched by m_A).

The proof of each case for m_C follows the same case in the proof of Lemma 4.1, where $p < p_j$ takes the place of pub_j . We show two example cases: *c_msg2a* and *c_dy2*

Case *c_msg2a*. Peer $j \rightarrow DY : sig_{sk_j}(p_j, x)$. If $\#\{k \mid 1 \leq k \leq t \wedge sig_{sk_k}(p_j, x) \in E\} = 2t - n - 1$ and $\#\{k \mid 1 \leq k \leq t \wedge sig_{sk_k}(p_j, x) \in E'\} = 2t - n$ then matched by $m_A = a_msg1$ for x, p_j . Otherwise matched by *skip*.

Case c_dy2 . If $x \in ITEM$ and $sig_{SSK}(p, x) \notin E$ and $sig_{SSK}(p, x) \in E'$, then this is matched by a_msg2 . It remains to show that $x \in R_p$ and $sig_{SSK}(p, B) \in E_A \Rightarrow x \in B$. The proof follows that of the same case in Lemma 4.1.

If $B_C \subseteq ITEM$ and $sig_{SSK}(p, B_C) \notin E$ and $sig_{SSK}(p, B_C) \in E'$, then this is matched by a_msg3 , with $Y = B_C$. The proof that (1) $E_A \cap PUBLISH_p = \{\}$, (2) $C_p \subseteq Y$ and (3) $Y \subseteq R_p$, for $Y = B_C$ is entirely similar to this case in the proof of Lemma 4.1.

Otherwise matched by *skip*.

The other cases follow the same pattern.

This concludes the proof that $BBSpec2 \preceq BBProt2$ with respect to $\{post, ack, publish\}$.

6 Accepting and Rejecting Posts

We now augment the Bulletin Board with an additional feature required for our use with Prêt à Voter: the ability to reject posts if they conflict with posts already received. For example, different votes cannot be accepted on the same ballot, and audit requests cannot be accepted (even on different boards) after a vote has been cast.

In particular, the bulletin board may refuse posts if they are inconsistent with previously accepted posts. We express this by introducing an irreflexive symmetric binary relation *clash* such that $clash(x, x')$ captures when two items x and x' should not both appear on the bulletin board. For convenience we define $clashset(x) = \{x' \mid clash(x, x')\}$ to be the set of all events that clash with x .

We will require that x will be accepted if there is no x' already received on any of the bulletin boards which clashes with x : in other words, that $clashset(x) \cap (\bigcup_p R_p) = \{\}$.

For example, in our context if x is a vote on a ballot then $clashset(x)$ will be the set of audits and other votes on that ballot. If x is an audit on a ballot then $clashset(x)$ will be the set of all possible votes on that ballot. If c is a cancellation of a ballot then $clashset(x) = \emptyset$: a cancellation can always be added to the bulletin board.

6.1 Specification

The specification is obtained by simply strengthening the guard in $BBSpec2$ of the event a_msg1 to include the non-clashing requirement. All other events are identical to those in $BBSpec2$. This yields the machine $BBSpec3$ as follows:


```

machine BBSpec3
variables  $E_A, R_p, C_p \quad (p \in \mathbb{N})$ 
    :
     $a\_msg1 \hat{=}$ 
      any  $x, p$ 
      where  $x \in E_A \cap ITEM \wedge p \in \mathbb{N} \wedge clashset(x) \cap (\bigcup_p R_p) = \{\}$ 
      then  $R_p := R_p \cup \{x\}$ 
      end;
    :
end

```

6.2 Implementation

We already have that *WBBProt2* is already a refinement of *BBSpec2*. Hence to obtain a refinement of *BBSpec3* it is enough to strengthen the guards of the events in *WBBProt2* matched by *a_msg1*, to ensure that when they are enabled then so is *a_msg1*. In order to complete the refinement proof we also need to strengthen the invariant with clauses 15 and 16 below.

In fact the only event matched by *a_msg1* in the proof of Lemma 4.1 is *c_msg2a*. We will thus obtain machine *WBBProt3* from *WBBProt2* by strengthening *c_msg2a* as follows:

```

 $c\_msg2a_j \hat{=}$ 
  any  $x$ 
  where  $x \in I_{j,p_j} \wedge clashset(x) \cap \{y \mid sk_j(p, y) \in \bigcup_p D_{j,p}\} = \{\}$ 
  then  $E := E \cup \{sig_{sk_j}(p_j, x)\} \parallel D_{j,p_j} := D_{j,p_j} \cup \{sig_{sk_j}(p_j, x)\}$ 
  end;

```

We also require two more clauses in the invariant of *WBBProt3*. It is straightforward to establish that all of the events of *WBBProt3* preserve these additional clauses:

$$sig_{sk_j}(p, x) \in E \Leftrightarrow sig_{sk_j}(p, x) \in D_{j,p} \quad (15)$$

$$clash(x, x') \wedge sig_{sk_j}(p, x) \in D_{j,p} \Rightarrow sig_{sk_j}(p', x') \notin \bigcup_p D_{j,p} \quad (16)$$

6.3 Simulation

With the exception of *a_msg1* and *c_msg2a*, all events in *WBBSpec3* and *WBBProt3* are exactly the same as in *BBSpec2* and *WBBProt2*, and so the refinements established previously remain valid.

We therefore only one new case to consider: *c_msg2a*:

Case c_msg2a . If (1) $\#\{k \leq t \mid sig_{sk_k}(p_j, x) \in E\} = 2t - n - 1$ and $sig_{sk_j}(p_j, x) \notin E$ and $clashset(x) \cap \{y \mid sk_j(p, y) \in \bigcup_p D_{j,p}\} = \{\}$ then this move will be matched by a_msg1 . Otherwise (2) c_msg2a is matched by $skip$ and we are done.

Hence for (1) it remains to prove that the guard of a_msg1 is enabled in this case, i.e. that $clashset(x) \cap (\bigcup_p R_p) = \{\}$. Since $J(s_A, s_C)$ this means that we must prove that in state s_C there is no $x' \in clashset(x)$ such that $x' \in \bigcup_p R_p$, i.e. no x' such that $\#\{k \leq t \mid sig_{sk_k}(p', x') \in E\} \geq 2t - n$.


We establish this by contraction. Assume there is some x' such that $clash(x, x')$ and $\#\{k \leq t \mid sig_{sk_k}(p', x') \in E\} \geq 2t - n$ in state s_C . This will also be the case in s'_C . Also in state s'_C we have $\#\{k \leq t \mid sig_{sk_k}(p_j, x) \in E\} = 2t - n$. Hence by Corollary 3.2 there is some $k < t$ such that $sig_{sk_k}(p', x') \in E$ and $sig_{sk_k}(p_j, x) \in E$. Hence by invariant (15) we have $sig_{sk_k}(p', x') \in D_{k'}$ and $sig_{sk_k}(p_j, x) \in D_{k,p_j}$. This yields a contradiction with invariant (16), since $clash(x, x')$.

We thus conclude that the guard of a_msg1 is enabled, and the refinement follows.

It follows that $BBSpec3 \preceq BBProt3$ with respect to $\{post, ack, publish\}$.

6.4 Example: lower threshold allows acceptance of clashing posts

Example 6.1. This example provides a second illustration as to why the threshold of honest peers is required to be greater than $2n/3$. If it is not, then an adversary can arrange for receipts to be issued on clashing posts, as shown:



In this attack two conflicting posts, x and y , are both provided with receipts. This is possible because no honest peer is involved in both: Peer 1 and (dishonest) Peer 3 contribute to the receipt of x , and Peers 2 and 3 contribute to the receipt of y . However, the WBB should only accept at most one of x and y . The attack is possible because Peer 3 provides shares towards the receipts of both x and y , something no honest peer would do. If the threshold of honest peers is greater than $2n/3$ then the same attack would require an honest peer to accept both x and y , which the protocol prevents.

7 Optimistic Commitment and Fallback

The final element of the bulletin board to introduce is the optimistic protocol within the publication phase, and the use of signed hashes in publication of the bulletin board.

Publication

When the time comes to publish, then the Post and Acknowledge protocol stops, and the peer begins the commit protocol which is used for the peers to obtain agreement on the bulletin board to publish.

Earlier models have used a round of message exchanges where peers circulate their database D_j , and another round where they circulate part-signed copies of their version of the bulletin board $t(D_j)$. In the case where there is some disagreement on databases then peers can update their bulletin boards to include new items they have received.

In practice we hope that in most cases the peers will agree on their local databases, and in this case they do not need to circulate them. We therefore introduce an optimistic commit where they can simply circulate a partially signed hash of their bulletin board (together with the period p): if they agree on the hash then they combine to give a threshold signature, and any peer can publish the bulletin board with the signed hash. If they do not agree then they can fall back to circulating their databases.

We therefore replace event c_msg6 by two messages: one to circulate a part-signed hash of the bulletin board $sig_{ssk_j}(h(t(D_j)))$, and one to circulate the bulletin board $t(D_j)$ itself (since this cannot be retrieved from the hash). The reason for separating these into two events is that we will eventually wish to schedule them separately: circulation of the hash will happen in the optimistic round, whereas publication of the board itself need not occur until there is agreement on the hash.

Optimistic commit protocol: This consists of two rounds:

1. $P_i \rightarrow P_j$: $sig_{ssk_i}(p, h(B_{i,p}))$ (for each $i, j \in I, j \neq i$)
each P_i checks the hashes from all peers agree
2. $P_i \rightarrow WBB$: $B_{i,p}, sig_{ssk_i}(p, h(B_{i,p}))$ (for each $i \in I$)

If there are not a threshold number of matching messages then the Fallback commit protocol is run:

Fallback commit protocol: This consists of a round of communications in which the peers exchange their databases in order to make them consistent.

1. $P_i \rightarrow P_j$: $D_{i,p}$ (for each $i, j \in I, j \neq i$)

Peer j receives $D_{i,p}$ from each of the other peers, up to some timeout. For each D received, Peer j adds to $D_{j,p}$ any item $sig_{ssk_k}(p, x) \in D$ that is not already in $D_{j,p}$.

The peers then return to the optimistic commit protocol.

7.1 Specification

The specification *BBSpec4* is similar to *BBSpec3* except that the form of the published bulletin board is changed, so that the unsigned bulletin board is published together with a signed hash. Events a_msg2 and a_msg3 are updated to reflect the change to the form in which

bulletin boards are published, and *publish* is also updated to output the new form of bulletin board.

$$\begin{aligned}
RECEIPT4 &= RECEIPT2 \\
PUBLISH4 &= \{sig_{SSK}(p, h(Y)) \mid p \in \mathbb{N} \wedge Y \subseteq ITEM\} \\
PUBLISH4_p &= \{sig_{SSK}(p, h(Y)) \mid Y \subseteq ITEM\}
\end{aligned}$$

```

machine BBSpec4
variables  $E_A, R_p, C_p \quad (p \in \mathbb{N})$ 
invariant  $E_A \subseteq ITEM \cup RECEIPT4 \cup PUBLISH4$ 
            $R_p \subseteq ITEM \quad (p \in \mathbb{N})$ 
            $C_p \subseteq ITEM \quad (p \in \mathbb{N})$ 

events
  init  $\hat{=}$   $E_A := \{\} \parallel \parallel_{p \in \mathbb{N}} (R_p := \{\} \parallel C_p := \{\})$ 
  post( $x$ )  $\hat{=}$  when  $x \in ITEM$  then  $E_A := E_A \cup \{x\}$  end;
   $r \leftarrow \text{ack} \hat{=}$   $r := (E_A \cap RECEIPT4)$ ;
   $P \leftarrow \text{publish} \hat{=}$ 
    any  $Y, p$ 
    where  $Y \subseteq ITEM \wedge sig_{SSK}(p, h(Y)) \in E_A$ 
    then  $P := (Y, sig_{SSK}(p, h(Y)))$ 
    end;
  a_msg1  $\hat{=}$ 
    any  $x \in E_A \cap ITEM \wedge p \in \mathbb{N}$ 
    then  $R_p := R_p \cup \{x\}$ 
    end;
  a_msg2  $\hat{=}$ 
    any  $x, p$ 
    where  $x \in R_p \wedge (sig_{SSK}(p, h(B)) \in E_A \Rightarrow x \in B)$ 
    then  $E_A := E_A \cup \{sig_{SSK}(p, x)\} \parallel C_p := C_p \cup \{x\}$ 
    end;
  a_msg3  $\hat{=}$ 
    any  $Y, p$ 
    where  $C_p \subseteq Y \subseteq R_p \wedge E_A \cap PUBLISH4 = \{\}$ 
    then  $E_A := E_A \cup \{sig_{SSK}(p, h(Y))\}$ 
    end
end

```

7.2 Implementation

Shared signatures on the bulletin board are now on its hash value. The effects of this change is highlighted in the invariant below. The remaining clauses remain the same as in *BBProt3*.

```

machine BBProt4
refines BBSpec4
variables  $E, I_{j,p}, D_{j,p}, H_{j,p}, p_j \ (1 \leq j \leq t)$ 

invariant
/* Types */
 $p_j \in \mathbb{N}$ 
 $E \subseteq MESSAGE$ 
 $I_{j,p} \subseteq ITEM$ 
 $D_{j,p} \subseteq \{sig_{sk_k}(p, x) \mid x \in ITEM\}$ 
 $H_{j,p} \subseteq \{sig_{sk_k}(p, h(B)) \mid B \subseteq ITEM\}$ 
/* Key invariant properties */
 $k \leq t \wedge k \in c[p, x] \wedge k \in s[p, B] \Rightarrow x \in B$ 
 $sig_{ssk_j}(p, x) \in E \Rightarrow \#d_j[p, x] \geq t$ 
 $sig_{SSK}(p, x) \in E \Rightarrow \#c[p, x] \geq t$ 

$sig_{ssk_j}(p, h(B)) \in E \Rightarrow B \subseteq t(D_{j,p})$



$sig_{SSK}(p, h(B)) \in E \Rightarrow \#s[p, B] \geq t$

 $D_{j,p} \subseteq E$ 
 $k \leq t \wedge k \in s[p, B] \Rightarrow c_k > p$ 
 $k \leq t \wedge k \in s[p, B_1] \wedge B_1 \neq B_2 \Rightarrow k \notin s[p, B_2]$ 
/* linking invariant */
 $R_p = \{x \in ITEM \mid \#\{k \mid 1 \leq k \leq t \wedge sig_{sk_k}(p, x) \in E\} \geq 2t - n\}$ 
 $C_p = \{x \in ITEM \mid sig_{SSK}(p, x) \in E\}$ 
 $E_A = E \cap (ITEM \cup RECEIPT4 \cup PUBLISH4)$ 
where:
 $d_j[p, x] = \{k \mid sig_{sk_k}(p, x) \in D_j\}$       shares of part sigs on  $x$  received by Peer  $j$ 
 $c[p, x] = \{k \mid sig_{ssk_k}(p, x) \in E\}$       peers which have (part)signed the receipt on  $x$ 

$s[p, B] = \{k \mid sig_{ssk_k}(p, h(B)) \in E\}$

      peers which have part-signed bulletin board  $B$ 

```

External events

External events are very similar in *BBProt4*. The event *publish* is adjusted to reflect the new form of publication, but external events are otherwise the same as in *BBProt3*.

```

events
  init  $\hat{=}$   $E := \{sk_k \mid k > t\} \cup \{ssk_k \mid k > t\} \parallel$ 
     $\parallel_{j,n} (I_{j,n} := \{\} \parallel D_{j,n} := \{\} \parallel H_{j,n} := \{\} \parallel p_j := 0 \parallel c_j := 0);$ 
  post( $x$ )  $\hat{=}$  when  $x \in ITEM$  then  $E := E \cup \{x\}$  end;
   $r \leftarrow \text{ack} \hat{=}$   $r \in (E \cap RECEIPT4);$ 
   $P \leftarrow \text{publish} \hat{=}$ 
    any  $Y, p$ 
    where  $Y \subseteq ITEM \wedge Y \in E \wedge sig_{SSK}(p, h(Y)) \in E$ 
    then  $P := (Y, sig_{SSK}(p, h(Y)))$ 
    end;

```

Posting and acknowledgement protocol

Posting and acknowledgement is identical to *BBProt3*, and so we do not reproduce the events here.

Commit and publish protocol

Agreeing and publishing the bulletin board now has two additional events: sending the bulletin board explicitly in *c_msg7*, and circulating signed hashes in the optimistic phase *c_msg8*. Note that *c_msg6* now provides a partially signed hash rather than a partially signed bulletin board.

```

c_msg4j  $\hat{=}$       /* start commit protocol */
  begin
     $p_j := p_j + 1$ 
  end;

c_msg5aj  $\hat{=}$       /* send database */
  any  $p < p_j$ 
  then  $E := E \cup \{D_{j,p}\}$ 
  end;

c_msg5bj  $\hat{=}$       /* receive database */
  any  $D, p$ 
  where  $D \in E \wedge D \subseteq SIG1_p \wedge p < p_j$ 
  then  $D_{j,p} := D_{j,p} \cup D$ 
  end;

c_msg6j  $\hat{=}$       /* provide partially signed hash */
  when  $c_j < p_j \wedge \#\{k \mid sig_{sk_k}(c_j, h(t(D_{j,c_j}))) \in H_{j,c_j}\} \geq t$ 
  then  $E := E \cup \{sig_{ssk_j}(c_j, h(t(D_{j,c_j})))\} \parallel c_j := c_j + 1$ 
  end;

c_msg7j  $\hat{=}$       /* send bulletin board */
  any  $p < p_j$ 
  then  $E := E \cup \{t(D_{j,p})\}$ 
  end;

c_msg8aj  $\hat{=}$       /* send signed hash */
  any  $p < p_j$ 
  then  $E := E \cup \{sig_{sk_j}(p, h(t(D_{j,p})))\}$ 
  end;

c_msg8bj  $\hat{=}$       /* receive signed hash */
  any  $p < p_j \wedge k \leq n \wedge B \subseteq ITEM \wedge sig_{sk_k}(p, h(B)) \in E$ 
  then  $H_{j,p} := H_{j,p} \cup \{sig_{sk_k}(p, h(B))\}$ 
  end;

```

Dolev-Yao environment

The adversary has the same moves as before, with one additional one: hashing. Any message can be hashed.

```

c_dy8  $\hat{=}$       /* hashing */
  any  $m$ 
  where  $m \in E$ 
  then  $E := E \cup \{h(m)\}$ 
  end

```


7.3 Simulation

The proof of simulation follows exactly the pattern of previous proofs. In particular:

- The external events *post*, *ack* and *publish* of *BBSpec4* are refined by their counterparts in *BBSpec4*.
- Event *a_msg1* is refined by the appropriate occurrence of *c_msg2a*.
- Events *a_msg2* and *a_msg3* are refined by *c_dy2* combining the signature shares appropriate to each case.
- All other concrete events refine *skip*. In particular, the new events of the optimistic protocol *c_msg7* and *c_msg8* refine *skip*.

This concludes the proof that $BBSpec4 \preceq BBProt4$ with respect to $\{post, ack, publish\}$, establishing the correctness of the bulletin board protocol *BBProt4* against the specification *BBSpec4*.

8 Liveness

The Dolev-Yao threat model does not allow the protocol to provide any liveness guarantees. All communications between the parties involved in the protocols can be blocked, preventing protocols from completing. The threat model is appropriate for consideration of safety properties, but is too strong for analysis of liveness. We require some assumptions about the communications between the protocol participants, as well as their honesty, in order to consider liveness.

We are primarily concerned with liveness for the publication of the bulletin board at the end of each period. To reason about liveness we assume that communication between peers is reliable, but that some of the peers may not follow the protocol, either because they are dishonest, or because they have failed.

8.1 All honest peers

We consider the case where all bulletin board peers are honest and follow the protocol. This scenario includes the case where users may be dishonest, sending different information to different peers, or not involving peers in some posting rounds. It also allows for the possibility where peers have not synchronised perfectly on the end of the period, so some posted items may be received in different periods for different peers, and hence their local records of the bulletin boards will not match.

Liveness can be shown for the commit and publish protocol. Different peers may begin that protocol with different databases D_i from the postings in the period. The optimistic protocol may complete if enough of them agree on $t(D_i)$, the contents of the bulletin board. However, it might not complete if the peers have sufficiently different records of what the

bulletin board should contain. In that case all peers execute the fallback protocol, and so communicate their databases reliably to all other peers. This results in all peers ending up with the same database record of posted items (i.e. the union of all their databases), and thus the second execution of the optimistic protocol will succeed in generating a threshold signature on the bulletin board.

8.2 A threshold of honest peers, and honest users

We now consider the case where some peers are not in communication for the commit and publish protocol. However we assume a threshold are behaving correctly and communicating with one another.

If every post of an item during the posting phase involved a threshold of (not necessarily the same) peers and obtained a receipt, then a single round of the fallback phase will ensure that all posted items are now obtained by all of the live peers. They are all sharing their evidence, and for each post there is at least one honest peer who is live in the commit and publish phase and also participated in the receipt of that item. This peer will provide the evidence of receipt to the other peers in the fallback round. Hence they will all agree on all posted items in the subsequent optimistic round. In this case again only one fallback round is required before agreement was obtained.

8.3 A threshold of honest peers

We now consider the more general case, where only a threshold set of peers are honest and connected during the publication phase. All peers will fix on a database D_i when they enter the publication phase: honest peers will use the D_i corresponding to the item posts they have received in the period, and we allow that dishonest peers will select any arbitrary D_i within their capability. We assume that peers will not change their database once it has been fixed, and will not send different databases to different peers, since this form of dishonesty would be easy to detect in practice. If it is detected then the dishonest peer would be removed and the protocol re-run (corresponding to a complete failure of that peer). Thus the only failure we need consider for peers outside the threshold set is failure to communicate, known as a stopping failure.

With this form of failure the *Floodset* agreement algorithm [Lyn96, 6.2] will lead to all honest peers agreeing on a database, and hence a bulletin board, within a maximum of $n - t + 1$ rounds. Each round of the Floodset algorithm is essentially the fallback protocol, with the optimistic protocol checking after each round whether there is a consensus. In the context of our commit and publish protocol the peers are looking for agreement on the union of their values rather than on one particular value, and so take their agreed value to be that set.

9 Discussion

9.1 Summary

In this paper we have presented a distributed protocol for running a bulletin board using a number of peers, which can tolerate a number of them failing, in the context of a threat model which has the communication between the peers controlled by a Dolev-Yao adversary, who also is able to control some of the peers. This provides robustness and distributed trust: the bulletin board can tolerate some peers failing, and we require that a threshold of peers should be honest but no individual peer is required to be trusted. Provided a threshold of the peers behave according to the protocol, the key properties demanded of the bulletin board hold. In particular, only items posted to the bulletin board will be posted on it, any item whose receipt is acknowledged by the bulletin board must be posted on it, and the bulletin board will not accept conflicting items. The bulletin board protocol has also been shown to be live when a threshold of honest peers all communicate without interference with each other, even if some dishonest peers attempt to disrupt progress.

The development of this modelling and verification approach for this kind of protocol is also one of the contributions of this paper. Correctness has been established formally using the Event-B framework, using simulation to show that the protocol is a refinement of an idealised bulletin board which has the desired properties. The model included a Dolev-Yao attacker and the description of the protocol steps followed by the peers. Carrying out the proof identified some nondeterminism inherent in the protocol and enabled us to include it in the idealisation to document the possible behaviour of the implementation. In particular, an adversary can create a situation where he controls whether or not an unreceipted item appears on the bulletin board, and so this is reflected as nondeterminism at the abstract level. In the context of the vVote system this will not be an issue in practice since the voting ceremony requires that any unreceipted items should be cancelled. Hence the nondeterminism will not affect the tallying of the election: either the cancellation appears alongside a vote, or it appears without the vote.

9.2 Related work

Other proposals for bulletin board implementations using a set of peers are given by Krummenacher [Kru10], by Peters [Pet05], and in the STAR-VOTE system [BBB⁺13].

Krummenacher’s Bulletin Board

Krummenacher focuses on a peered bulletin board that guarantees the correctness of its history and the authenticity of the messages. His proposal is motivated by the desire to provide a distributed version of Heather and Lundin’s append-only web bulletin board [HL08]. The protocol is designed essentially for robustness, and is considered in the context where up to k out of n peers may fail. A particular number of peers ($k + 1$) must accept an item for it to be allowed onto the bulletin board. Peers hold their own versions of the history of items posted, and so their histories would need to be combined in order to obtain the global bulletin

board. A similarity with our approach is the need for peers to confirm that other peers have received an item before providing their own response. However, peers use a locking protocol when they seek confirmation from other peers, so the approach does not scale up well as k gets larger relative to n .

The most significant difference with the approach of this paper is the threat model: Krummenacher considers the protocol in the context of communication failures and peer failures, but does not consider an active adversary or corrupt peers who might deliberately introduce invalid messages (or accept clashing items). Formal modelling and analysis would help to clarify the adversarial context and identify whether the protocol does indeed guarantee correct behaviour within that threat model. Another difference is that Krummenacher’s bulletin board is not concerned with preventing clashing items from being posted. This might be addressed by setting the threshold k to be sufficiently high and requiring that individual peers do not accept items that clash. If peers can be dishonest then we may require the threshold of $k > 2n/3$, but this threshold does not work well with the locking protocol used in posting. Finally, we also observe that the protocol follows the approach of [HL08] in using timestamps to ensure that the bulletin board is relatively recent, and hence that no commit round is required. Instead the peers are always able to provide their current version of the bulletin board. This approach gives rise to challenges in implementation, notably that a single view of the ‘official’ bulletin board, as would be required in an election context, would need to be constantly refreshed by the bulletin board peers. This would be a substantial overhead, and furthermore its security implications are not well understood. For all these reasons Krummenacher’s implementation is not suitable as it stands for our requirements.

Peters’ Bulletin Board

Peters [Pet05] considers several approaches, and proposes and implements a bulletin board which uses a secure agreement protocol [Rei94] on top of a group membership protocol [Rei96]. Items are posted by a client to a single peer, which then communicates with the others, obtains confirmation of receipt from a threshold of them, circulates that confirmation back to the peers, and returns a receipt to the client. This approach is similar to our posting protocol, where peers require confirmation of receipt from a threshold of other peers, before returning their share of the signature on the receipt. The system requires the same threshold as we do: that strictly more than $2/3$ of the peers behave honestly. Further, each honest peer can serve the complete bulletin board on request. They achieve this by means of a round of communication collecting signature shares on the bulletin board after each item is posted, similar to our optimistic protocol for the end-of-period publication. In practice this might carry an overhead, both in obtaining the agreement and in providing the bulletin board, and in our context it is not necessary. However, it would be perfectly possible to run the bulletin board and only carry this out at the end of the period.

Similar to our approach, peers can also reject posts that clash with previous posts (such as a second vote on the same ballot form), and the threshold ensures that the collective bulletin board will not accept posts that clash. Dishonest peers are handled by use of the group membership protocol: a current group of participating peers is maintained by all honest peers, and dishonest peers once detected are removed from the group. Peers can also be

readmitted to the group, in which case they need to bring themselves up to date on the state of the bulletin board.

A key difference with our approach is the use of the group membership protocol to dynamically change the set of ‘live’ peers. Peters provides excellent formal descriptions of the protocols with sufficient detail to enable code production, and also gives arguments of their correctness in this setting. However, there is no formal verification of the collection of protocols operating together. The possible interactions between them are quite subtle and require careful handling, for example how reconfiguration of the group might interact with the posting of items, or how a client may need to switch from an ejected peer to an honest one. Some dishonest peer behaviour might not trigger removal from the group, but might still interfere with the protocols, and this possibility requires careful analysis. A second key difference is the way the (honest) peers need to keep a record of the up to date bulletin board at all times. Although honest peers in our system will also have a record of the bulletin board if they are connected and participate in the posting of items, it is not a requirement, and peers are not relied on for it. A final key difference is in the threat model, which allows dishonest peers but considers the network itself to be reliable (as we do for liveness), so honest peers can always communicate with each other. A peer failing to communicate is treated as dishonest. This will trigger a reconfiguration of the group and makes the protocol more sensitive to minor communication failures.

STAR-Vote

The use of the bulletin board within STAR-Vote is close to ours: it collects votes during the election, and publishes only at the end. The voting terminals are networked and play the role of bulletin board peers: they collect the votes as they are cast, and track which ones are to be counted. The voter retains a paper receipt as evidence of their vote, which does not reveal anything of how they voted. The system is designed to tolerate faulty peers. At the end of the election the voting terminals agree on the votes that have been received, and publish them in encrypted form on the web so that voters can check them against their receipts. The electronic record is also checked against the paper copies of the votes retained by the system. The set of votes is signed by the election authorities to prevent subsequent manipulation. Following the approach of VoteBox[SDW08], the Voting Terminals maintain a global audit log during the election using a hash chain, so received votes are committed to in real time, and past events cannot be tampered with by a subset of malicious machines.

The main difference with our approach is that STAR-Vote is designed for use in a single polling station. This gives a different threat model, in particular the threats are considered to come from malicious devices rather than the underlying network. There is no geographic separation between casting a vote and having it received by the bulletin board peers, and the local network is assumed to provide reliable communication. STAR-Vote therefore does not need to address the challenge of posting items to a remote bulletin board, which we have had to address by having the posting protocol generate a cryptographically signed receipt to provide the evidence that the item has been received. Currently STAR-Vote does not provide signatures on the receipts that voters retain, though this is considered as a possibility in the context of mitigation against a “defaming” attack where voters present falsified evidence against the bulletin board. STAR-Vote also does not go into detail about how the bulletin

board information is collated from the peers, in particular what happens when some voting terminals but not others claim to have received a vote. The emphasis is on detection of incorrect behaviour rather than its automatic correction. If discrepancies are identified, then the approach would be to resolve them forensically, by checking audit logs, memory dumps, and other relevant records.

Byzantine Agreement Protocols

The general problem of achieving generalised agreement across components where some might fail in adverse ways is known as the *Byzantine Agreement* problem [LSP82], and there is an extensive literature on approaches to the problem [Lyn96]. Such protocols require correct behaviour in strictly more than $2/3$ of the peers, the same requirement as we have on our bulletin board peers. However, Byzantine Agreement protocols are not really suitable for items being posted. These protocols are typically synchronous and proceed in rounds, which would be too inefficient for receiving large quantities of votes: too many rounds and too much synchronisation overhead would be required to process each vote if we wish the peers to agree on the receipt of every vote. Furthermore, not all peers would necessarily be aware when a protocol run is starting, since they may not receive the initial item. Instead we have provided a protocol for the peers simply to send messages to each other and to respond to messages received in an asynchronous fashion. This means that the peers do not all need to agree on each vote. Our threat model is also different to the typical threat model for Byzantine agreement protocols: ours allows honest peers to be excluded from the acceptance of some items to the bulletin board, without being considered as dishonest, whereas Byzantine agreement protocols consider any non-participating peer as failing.

We are closer to the problem of Byzantine agreement in the commit and publish phase, since this is where the peers seek consensus to agree on a bulletin board to publish. Our optimistic and fallback protocols are indeed close to the Floodset protocol [Lyn96], a basic agreement protocol. Even in this case we do not require the full power of Byzantine agreement protocols: the use of signatures minimises the ability of dishonest peers to introduce additional confusing information to disrupt the protocol run, and we can limit the adversarial behaviour simply to peers ceasing to communicate.

9.3 Implementation level considerations

The concrete model above has been analysed for correctness, and shown to be correct with respect to the abstract model. However, even the concrete model is nondeterministic in the order in which events should be scheduled. This is deliberate, since it means that interactions between protocols are addressed in the analysis, but in practice we will want an efficient implementation and so will schedule events in a particular way, to avoid expensive computations such as the fallback protocol unless they are necessary, or for other pragmatic reasons. For example, the implementation we have developed for the vVote project requires that all peers should agree on the hash of the bulletin board in order to provide their signature share, although only a threshold of hashes on the database is sufficient for correctness. We do this because it is still helpful to know if possible that all peers have the same database, to provide

reassurance that the protocol is working properly. However, this implementation is consistent with the concrete model, in which peers can start the fallback protocol at any time. As long as the implementation performs events in accordance with the concrete protocol, it will provide behaviour that is correct with respect to the abstract specification.

Once published, a web bulletin board will need provide voters with the facilities to confirm their vote is correctly recorded within the signed bulletin board, and to be able to obtain the full contents of the signed bulletin board so that the subsequent processing can be checked. In order to ensure that the board cannot be later replaced with a different signed bulletin board, the signed hash of the board will also be published at the end of the period using an out-of-band broadcast channel. For the planned use of the system in Victoria in November 2014, each period will be one day, and the signed hash will be published the following day in the newspaper. Voters can then check the bulletin board on the web against that published hash.

There are different ways to make the contents of the bulletin board available. To be consistent with the commit and publish protocol, all that is needed is the peers are able to produce shares on a threshold signature of what is produced. It may simply make the entire bulletin board available for download, and then the voter will check its signature and that their vote is included (and not cancelled) within it. Alternatively it may make use of a structure such as a hash tree [GTT09] which provides signature authentication that an item is included on the bulletin board without the need to download the entire board. In our case we need to check not only that a vote is present, but also that there is also no cancellation present. In practice this may be done by also publishing and signing the entire set of cancellations, since this is likely to be small compared to the set of votes cast. Checking the presence of a vote then involves confirming both that a vote is present in the hash tree and that it has not been cancelled. This is the subject of ongoing research.

Acknowledgements

We are grateful to James Heather, Peter Ryan, Vanessa Teague and Douglas Wikström for useful discussions on Bulletin Board design and approaches to verification. We are also grateful to Gavin Lowe and Joshua Guttman for detailed discussions on the formal modelling and verification approach, and to Thierry Lecomte, Helen Treharne, John Derrick and Graeme Smith for discussion and advice on the B modelling and refinement. Thanks also to Olivier Pereira and Dan Wallach for clarifying aspects of STAR-Vote. This work was supported by the EPSRC Trustworthy Voting Systems project EP/G025797/1.

References

- [Abr10] Jean-Raymond Abrial. *Modeling in Event-B - System and Software Engineering*. Cambridge University Press, 2010.
- [Adi06] Ben Adida. *Advances in Cryptographic Voting Systems*. PhD thesis, MIT Cambridge, July 2006.

- [Adi08] Ben Adida. Helios: Web-based open-audit voting. In *USENIX Security Symposium*, pages 335–348, 2008.
- [BBB⁺13] Susan Bell, Josh Benaloh, Michael D. Byrne, Dana DeBeauvoir, Bryce Eakin, Gail Fisher, Philip Kortum, Neal McBurnett, Julian Montoya, Michelle Parker, Olivier Pereira, Philip B. Stark, Dan S. Wallach, and Michael Winn. STAR-vote: A secure, transparent, auditable, and reliable voting star-vote: A secure, transparent, auditable, and reliable voting system. *USENIX Journal of Election Technology and Systems (JETS)*, 1(1), August 2013.
- [BCH⁺12] Craig Burton, Chris Culnane, James Heather, Thea Peacock, Peter Y. A. Ryan, Steve Schneider, Sriramkrishnan Srinivasan, Vanessa Teague, Roland Wen, and Zhe Xia. A supervised verifiable voting protocol for the Victorian Electoral Commission. In *Proc. 5th International Conference on Electronic Voting*, 2012.
- [BNFL⁺12] Jonathan Ben-Nun, Niko Fahri, Morgan Llewellyn, Ben Riva, Alon Rosen, Amnon Ta-Shma, and Douglas Wikström. A new implementation of a dual (paper and cryptographic) voting system. In *5th International Conference on Electronic Voting (EVOTE)*, 2012.
- [CCC⁺10] Richard Carback, David Chaum, Jeremy Clark, John Conway, Aleksander Essex, Paul S. Herrnson, Travis Mayberry, Stefan Popoveniuc, Ronald L. Rivest, Emily Shen, Alan T. Sherman, and Poorvi L. Vora. Scantegrity II municipal election at Takoma Park: The first e2e binding governmental election with ballot privacy. In *Proc. USENIX Security*, 2010.
- [CCM08] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a secure voting system. In *IEEE Symposium on Security and Privacy*, pages 354–368, 2008.
- [CGS97] Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A secure and optimally efficient multi-authority election scheme. In *Advances in Cryptology*, number 1233 in Lecture Notes in Computer Science, pages 103–118. Springer-Verlag, 1997.
- [CRS05] D. Chaum, P.Y.A. Ryan, and S. Schneider. A practical, voter-verifiable election scheme. In *European Symposium on Research in Computer Security*, number 3679 in Lecture Notes in Computer Science. Springer-Verlag, 2005.
- [DB01] John Derrick and Eerke Boiten. *Refinement in Z and Object-Z*. Springer, 2001.
- [DY83] Danny Dolev and Andrew Chi-Chih Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–207, 1983.
- [GTT09] Michael T. Goodrich, Roberto Tamassia, and Nikos Triandopoulos. Efficient authenticated data structures for graph connectivity and geometric search problems. *CoRR*, abs/0908.4116, 2009.
- [HL08] James Heather and David Lundin. The append-only web bulletin board. In *Formal Aspects in Security and Trust*, pages 242–256, 2008.

- [Kru10] Roland Krummenacher. Implementation of a web bulletin board for e-voting applications. In *MSE Seminar on E-Voting*. Institute for Internet Technologies and Applications, 2010.
- [KTV12] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Clash attacks on the verifiability of e-voting systems. In *IEEE Symposium on Security and Privacy*, pages 395–409, 2012.
- [LSP82] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [Lyn96] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [MAV05] Christophe Métayer, Jean-Raymond Abrial, and Laurent Voisin. Event-B language, 2005. RODIN Project Deliverable 3.2, <http://rodin.cs.ncl.ac.uk/deliverables/D7.pdf>, accessed 25/5/10.
- [Nor13] Norwegian Ministry of Local Government and Regional Development. VALG: The e-vote trial, 2013.
- [Pet05] R.A. Peters. A Secure Bulletin Board. Master’s thesis, Technische Universiteit Eindhoven, 2005.
- [RBH⁺09] Peter Y. A. Ryan, David Bismark, James Heather, Steve Schneider, and Zhe Xia. Prêt à Voter: a voter-verifiable voting system. *IEEE Transactions on Information Forensics and Security*, 4(4):662–673, 2009.
- [Rei94] Michael K. Reiter. Secure agreement protocols: Reliable and atomic group multicast in rampart. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, and Ravi S. Sandhu, editors, *ACM Conference on Computer and Communications Security*, pages 68–80. ACM, 1994.
- [Rei96] Michael K. Reiter. A secure group membership protocol. *IEEE Trans. Software Eng.*, 22(1):31–42, 1996.
- [SDW08] Daniel R. Sandler, Kyle Derr, and Dan S. Wallach. Votebox: A tamper-evident, verifiable electronic voting system. In *Proc. 17th USENIX Security Symposium*, 2008.
- [Sur13] University of Surrey. Software design for VEC vVote system V1.0, 2013.
- [Wag13] David Wagner. Remote voting: What can we do? Keynote Address, Electronic Voting Technology Workshop/Workshop on Trustworthy Elections, 2013.